arXiv:1105.1186v1 [cs.RO] 5 May 2011

Sampling-based Algorithms for Optimal Motion Planning

Sertac Karaman Emilio Frazzoli*

Abstract

During the last decade, sampling-based path planning algorithms, such as Probabilistic RoadMaps (PRM) and Rapidly-exploring Random Trees (RRT), have been shown to work well in practice and possess theoretical guarantees such as probabilistic completeness. However, little effort has been devoted to the formal analysis of the quality of the solution returned by such algorithms, e.g., as a function of the number of samples. The purpose of this paper is to fill this gap, by rigorously analyzing the asymptotic behavior of the cost of the solution returned by stochastic sampling-based algorithms as the number of samples increases. A number of negative results are provided, characterizing existing algorithms, e.g., showing that, under mild technical conditions, the cost of the solution returned by broadly used sampling-based algorithms converges almost surely to a non-optimal value. The main contribution of the paper is the introduction of new algorithms, namely, PRM* and RRT*, which are provably asymptotically optimal, i.e., such that the cost of the returned solution converges almost surely to the optimum. Moreover, it is shown that the computational complexity of the new algorithms is within a constant factor of that of their probabilistically complete (but not asymptotically optimal) counterparts. The analysis in this paper hinges on novel connections between stochastic sampling-based path planning algorithms and the theory of random geometric graphs.

Keywords: Motion planning, optimal path planning, sampling-based algorithms, random geometric graphs.

1 Introduction

The robotic motion planning problem has received a considerable amount of attention, especially over the last decade, as robots started becoming a vital part of modern industry as well as our daily life (Latombe, 1991; LaValle, 2006; Choset et al., 2005). Even though modern robots may possess significant differences in sensing, actuation, size, workspace, application, etc., the problem of navigating through a complex environment is embedded and essential in almost all robotics applications. Moreover, this problem is relevant to other disciplines such as verification, computational biology, and computer animation (Latombe, 1999; Bhatia and Frazzoli, 2004; Branicky et al., 2006; Cortes et al., 2007; Liu and Badler, 2003; Finn and Kavraki, 1999).

Informally speaking, given a robot with a description of its dynamics, a description of the environment, an initial state, and a set of goal states, the motion planning problem is to find a sequence of control inputs so as the drive the robot from its initial state to one of the goal states while obeying the rules of the environment, e.g., not colliding with the surrounding obstacles. An algorithm to address this problem is said to be *complete* if it terminates in finite time, returning a valid solution if one exists, and failure otherwise.

Unfortunately, the problem is known to be very hard from the computational point of view. For example, a basic version of the motion planning problem, called the generalized piano movers

---

*The authors are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA.

problem, is PSPACE-hard (Reif, 1979). In fact, while complete planning algorithms exist (see, e.g., Lozano-Perez and Wesley, 1979; Schwartz and Sharir, 1983; Canny, 1988), their complexity makes them unsuitable for practical applications.

Practical planners came around with the development of cell decomposition methods (Brooks and Lozano-Perez, 1983) and potential fields (Khatib, 1986). These approaches, if properly implemented, relaxed the completeness requirement to, for instance, *resolution completeness*, i.e., the ability to return a valid solution, if one exists, if the resolution parameter of the algorithm is set fine enough. These planners demonstrated remarkable performance in accomplishing various tasks in complex environments within reasonable time bounds (Ge and Cui, 2002). However, their practical applications were mostly limited to state spaces with up to five dimensions, since decomposition-based methods suffered from large number of cells, and potential field methods from local minima (Koren and Borenstein, 1991). Important contributions towards broader applicability of these methods include navigation functions (Rimon and Koditschek, 1992) and randomization (Barraquand and Latombe, 1993).

The above methods rely on an explicit representation of the obstacles in the configuration space, which is used directly to construct a solution. This may result in an excessive computational burden in high dimensions, and in environments described by a large number of obstacles. Avoiding such a representation is the main underlying idea leading to the development of sampling-based algorithms (Kavraki and Latombe, 1994; Kavraki et al., 1996; LaValle and Kuffner, 2001). See Lindemann and LaValle (2005) for a historical perspective. These algorithms proved to be very effective for motion planning in high-dimensional spaces, and attracted significant attention over the last decade, including very recent work (see, e.g., Prentice and Roy, 2009; Tedrake et al., 2010; Luders et al., 2010; Berenson et al., 2008; Yershova and LaValle, 2008; Stilman et al., 2007; Koyuncu et al., 2010). Instead of using an explicit representation of the environment, sampling-based algorithms rely on a collision checking module, providing information about feasibility of candidate trajectories, and connect a set of points sampled from the obstacle-free space in order to build a graph (roadmap) of feasible trajectories. The roadmap is then used to construct the solution to the original motion-planning problem.

Informally speaking, sampling-based methods provide large amounts of computational savings by avoiding explicit construction of obstacles in the state space, as opposed to most complete motion planning algorithms. Even though these algorithms are not complete, they provide *probabilistic completeness* guarantees in the sense that the probability that the planner fails to return a solution, if one exists, decays to zero as the number of samples approaches infinity (Barraquand et al., 1997) (see also Hsu et al., 1997; Kavraki et al., 1998; Ladd and Kavraki, 2004). Moreover, the rate of decay of the probability of failure is exponential, under the assumption that the environment has good "visibility" properties (Barraquand et al., 1997). More recently, the empirical success of sampling-based algorithms was argued to be strongly tied to the hypothesis that most practical robotic applications, even though involving robots with many degrees of freedom, feature environments with such good visibility properties (Hsu et al., 2006).

1.1 Sampling-Based Algorithms

Arguably, the most influential sampling-based motion planning algorithms to date include Probabilistic RoadMaps (PRMs) (Kavraki et al., 1996, 1998) and Rapidly-exploring Random Trees (RRTs) (Kuffner and LaValle, 2000; LaValle and Kuffner, 2001; LaValle, 2006). Even though the idea of connecting points sampled randomly from the state space is essential in both approaches, these two algorithms differ in the way that they construct a graph connecting these points.

The PRM algorithm and its variants are multiple-query methods that first construct a graph (the roadmap), which represents a rich set of collision-free trajectories, and then answer queries by computing a shortest path that connects the initial state with a final state through the roadmap. The PRM algorithm has been reported to perform well in high-dimensional state spaces (Kavraki et al., 1996). Furthermore, the PRM algorithm is probabilistically complete, and such that the probability of failure decays to zero exponentially with the number of samples used in the construction of the roadmap (Kavraki et al., 1998). During the last two decades, the PRM algorithm has been a focus of robotics research: several improvements were suggested by many authors and the reasons to why it performs well in many practical cases were better understood (see, e.g., Branicky et al., 2001; Hsu et al., 2006; Ladd and Kavraki, 2004, for some examples).

Even though multiple-query methods are valuable in highly structured environments, such as factory floors, most online planning problems do not require multiple queries, since, for instance, the robot moves from one environment to another, or the environment is not known a priori. Moreover, in some applications, computing a roadmap a priori may be computationally challenging or even infeasible. Tailored mainly for these applications, incremental sampling-based planning algorithms such as RRTs have emerged as an online, single-query counterpart to PRMs (see, e.g., Kuffner and LaValle, 2000; Hsu et al., 2002). The incremental nature of these algorithms avoids the necessity to set the number of samples a priori, and returns a solution as soon as the set of trajectories built by the algorithm is rich enough, enabling on-line implementations. Moreover, tree-based planners do not require connecting two states exactly and more easily handle systems with differential constraints. The RRT algorithm has been shown to be probabilistically complete (Kuffner and LaValle, 2000), with an exponential rate of decay for the probability of failure (Frazzoli et al., 2002). The basic version of the RRT algorithm has been extended in several directions, and found many applications in the robotics domain and elsewhere (see, for instance, Frazzoli et al., 2002; Bhatia and Frazzoli, 2004; Cortes et al., 2007; Branicky et al., 2006, 2003; Zucker et al., 2007). In particular, RRTs have been shown to work effectively for systems with differential constraints and nonlinear dynamics (LaValle and Kuffner, 2001; Frazzoli et al., 2002) as well as purely discrete or hybrid systems (Branicky et al., 2003). Moreover, the RRT algorithm was demonstrated in major robotics events on various experimental robotic platforms (Bruce and Veloso, 2003; Kuwata et al., 2009; Teller et al., 2010; Shkolnik et al., 2011; Kuffner et al., 2002).

Other sampling-based planners of note include Expansive Space Trees (EST) (Hsu et al., 1997, 1999) and Sampling-based Roadmap of Trees (SRT) (Plaku et al., 2005). The latter combines the main features of multiple-query algorithms such as PRM with those of single-query algorithms such as RRT and EST.

1.2 Optimal Motion Planning

In most applications, the quality of the solution returned by a motion planning algorithm is important. For example, one may be interested in solution paths of minimum cost, with respect to a given cost functional, such as the length of a path, or the time required to execute it. The problem of computing optimal motion plans has been proven in Canny and Reif (1987) to be very challenging even in basic cases.

In the context of sampling-based motion planning algorithms, the importance of computing optimal solutions has been pointed out in early seminal papers (LaValle and Kuffner, 2001). However, optimality properties of sampling-based motion planning algorithms have not been systematically investigated, and most of the relevant work relies on heuristics. For example, in many field implementations of sampling-based planning algorithms (see, e.g., Kuwata et al., 2009), it is often the case that since a feasible path is found quickly, additional available computation time is devoted to improving the solution with heuristics until the solution is executed. Urmson and Simmons

(2003) proposed heuristics to bias the tree growth in RRT towards those regions that result in low-cost solutions. They have also shown experimental results evaluating the performance of different heuristics in terms of the quality of the solution returned. Ferguson and Stentz (2006) considered running the RRT algorithm multiple times in order to progressively improve the quality of the solution. They showed that each run of the algorithm results in a path with smaller cost, even though the procedure is not guaranteed to converge to an optimal solution. Criteria for restarting multiple RRT runs, in a different context, were also proposed in Wedge and Branicky (2008). A more recent approach is the transition-based RRT (T-RRT) designed to combine rapid exploration properties of the RRT with stochastic global optimization methods (Jaillet et al., 2010; Berenson et al., 2011).

A different approach that also offers optimality guarantees is based on graph search algorithms, such as A^*, applied over a finite discretization (based, e.g., on a grid, or a cell decomposition of the configuration space) that is generated offline. Recently, these algorithms received a large amount of attention. In particular, they were extended to run in an anytime fashion (Likhachev et al., 2004, 2008), deal with dynamic environments (Stentz, 1995; Likhachev et al., 2008), and handle systems with differential constraints (Likhachev and Ferguson, 2009). These have also been successfully demonstrated on various robotic platforms (Likhachev and Ferguson, 2009; Dolgov et al., 2009). However, optimality guarantees of these algorithms are only ensured up to the grid resolution. Moreover, since the number of grid points grows exponentially with the dimensionality of the state space, so does the (worst-case) running time of these algorithms.

1.3 Statement of Contributions

To the best of the author's knowledge, this paper provides the first systematic and thorough analysis of optimality and complexity properties of the major paradigms for sampling-based path planning algorithms, for multiple- or single-query applications, and introduces the first algorithms that are both asymptotically optimal and computationally efficient, with respect to other algorithms in this class. A summary of the contributions can be found below, and is shown in Table 1.

As a first set of results, it is proven that the standard PRM and RRT algorithms are not asymptotically optimal, and that the "simplified" PRM algorithm is asymptotically optimal, but computationally expensive. Moreover, it is shown that the k-nearest variant of the (simplified) PRM algorithm is not necessarily probabilistically complete (e.g., it is not probabilistically complete for $k = 1$), and is not asymptotically optimal for any fixed k.

In order to address the limitations of sampling-based path planning algorithms available in the literature, new algorithms are proposed, i.e., PRM*, RRG, and RRT*, and proven to be probabilistically complete, asymptotically optimal, and computationally efficient. Of these, PRM* is a batch variable-radius PRM, applicable to multiple-query problems, in which the radius is scaled with the number of samples in a way that provably ensures both asymptotic optimality and computational efficiency. RRG is an incremental algorithm that builds a connected roadmap, providing similar performance to PRM* in a single-query setting, and in an anytime fashion (i.e., a first solution is provided quickly, and monotonically improved if more computation time is available). The RRT* algorithm is a variant of RRG that incrementally builds a tree, providing anytime solutions, provably converging to an optimal solution, with minimal computational and memory requirements.

In this paper, the problem of planning a path through a connected bounded subset of a d-dimensional Euclidean space is considered. As in the early seminal papers on incremental sampling-based motion planning algorithms such as Kuffner and LaValle (2000), no differential constraints are considered (i.e., the focus of the paper is on path planning problems), but our methods can be easily extended to planning in configuration spaces and applied to several practical problems of interest.

Table 1: Summary of results. Time and space complexity are expressed as a function of the number of samples n, for a fixed environment.

	Algorithm	Probabilistic Completeness	Asymptotic Optimality	Monotone Convergence	Time Complexity		Space Complexity
					Processing	Query	
Existing Algorithms	PRM	Yes	No	Yes	$O(n \log n)$	$O(n \log n)$	$O(n)$
	sPRM	Yes	Yes	Yes	$O(n^2)$	$O(n^2)$	$O(n^2)$
	k-sPRM	Conditional	No	No	$O(n \log n)$	$O(n \log n)$	$O(n)$
	RRT	Yes	No	Yes	$O(n \log n)$	$O(n)$	$O(n)$
Proposed Algorithms	PRM*	Yes	Yes	No	$O(n \log n)$	$O(n \log n)$	$O(n \log n)$
	k-PRM*						
	RRG	Yes	Yes	Yes	$O(n \log n)$	$O(n \log n)$	$O(n \log n)$
	k-RRG						
	RRT*	Yes	Yes	Yes	$O(n \log n)$	$O(n)$	$O(n)$
	k-RRT*						

The extension to systems with differential constraints is deferred to future work (see Karaman and Frazzoli (2010a) for preliminary results).

Finally, the results presented in this article, and the techniques used in the analysis of the algorithms, hinge on novel connections established between sampling-based path planning algorithms in robotics and the theory of random geometric graphs, which may be of independent interest.

A preliminary version of this article has appeared in Karaman and Frazzoli (2010b). Since then a variety of new algorithms based on the the ideas behind PRM*, RRG, and RRT* have been proposed in the literature. For instance, a probabilistically complete and probabilistically sound algorithm for solving a class of differential games has appeared in Karaman and Frazzoli (2010c). Algorithms based on the RRG were used to solve belief-space planning problems in Bry and Roy (2011). The RRT* algorithm was used for anytime motion planning in Karaman et al. (2011), where it was also demonstrated experimentally on a full-size robotic fork truck. In Alterovitz et al. (2011), the analysis given in Karaman and Frazzoli (2010b) was used to guarantee computational efficiency and asymptotic optimality of a new algorithm that can trade off between exploration and optimality during planning.

A software library implementing the new algorithms introduced in this paper has been released as open-source software by the authors, and is currently available at `http://ares.lids.mit.edu/software/`

1.4 Paper Organization

This paper is organized as follows. Section 2 lays the ground in terms of notation and problem formulation. Section 3 is devoted to the discussion of the algorithms that are considered in the paper: first, the main paradigms for sampling-based motion planning algorithms available in the literature are presented, together with their main variants. Then, the new proposed algorithms are presented and motivated. In Section 4 the properties of these algorithms are rigorously analyzed, formally establishing their probabilistic completeness and asymptotically optimality (or lack thereof), as well as their computational complexity as a function of the number of samples and of the number of obstacles in the environment. Experimental results are presented in Section 5, to illustrate and validate the theoretical findings. Finally, Section 6 contains conclusions and perspectives for future work. In order not to excessively disrupt the flow of the presentation, a summary of notation used throughout the paper, as well as lengthy proofs of important results are presented in the Appendix.

2 Preliminary Material

This section contains some preliminary material that will be necessary for the discussion in the remainder of the paper. Namely, the problems of feasible and optimal motion planning is introduced, and some important results from the theory of random geometric graphs are summarized. The notation used in the paper is summarized in Appendix A.

2.1 Problem Formulation

In this section, the feasible and optimal path planning problems are formalized.

Let $\mathcal{X} = (0,1)^d$ be the *configuration space*, where $d \in \mathbb{N}$, $d \geq 2$. Let $\mathcal{X}_{\text{obs}}$ be the *obstacle region*, such that $\mathcal{X} \setminus \mathcal{X}_{\text{obs}}$ is an open set, and denote the *obstacle-free space* as $\mathcal{X}_{\text{free}} = \text{cl}(\mathcal{X} \setminus \mathcal{X}_{\text{obs}})$, where $\text{cl}(\cdot)$ denotes the closure of a set. The *initial condition* x_{init} is an element of $\mathcal{X}_{\text{free}}$, and the *goal region* $\mathcal{X}_{\text{goal}}$ is an open subset of $\mathcal{X}_{\text{free}}$. A path planning problem is defined by a triplet $(\mathcal{X}_{\text{free}}, x_{\text{init}}, \mathcal{X}_{\text{goal}})$.

Let $\sigma : [0,1] \to \mathbb{R}^d$; the *total variation* of σ is defined as

$$\text{TV}(\sigma) = \sup_{\{n \in \mathbb{N}, 0=\tau_0<\tau_1<\cdots<\tau_n=s\}} \sum_{i=1}^{n} |\sigma(\tau_i) - \sigma(\tau_{i-1})|.$$

A function σ with $\text{TV}(\sigma) < \infty$ is said to have *bounded variation.*

Definition 1 (Path) *A function $\sigma : [0,1] \to \mathbb{R}^d$ of bounded variation is called a*

- Path*, if it is continuous;*
- Collision-free path*, if it is a path, and $\sigma(\tau) \in \mathcal{X}_{\text{free}}$, for all $\tau \in [0,1]$;*
- Feasible path*, if it is a collision-free path, $\sigma(0) = x_{\text{init}}$, and $\sigma(1) \in \text{cl}(\mathcal{X}_{\text{goal}})$.*

The total variation of a path is essentially its length, i.e., the Euclidean distance traversed by the path in $\mathbb{R}^d$. The *feasibility problem* of path planning is to find a feasible path, if one exists, and report failure otherwise:

Problem 2 (Feasible path planning) *Given a path planning problem $(\mathcal{X}_{\text{free}}, x_{\text{init}}, \mathcal{X}_{\text{goal}})$, find a feasible path $\sigma : [0,1] \to \mathcal{X}_{\text{free}}$ such that $\sigma(0) = x_{\text{init}}$ and $\sigma(1) \in \text{cl}(\mathcal{X}_{\text{goal}})$, if one exists. If no such path exists, report failure.*

Let Σ denote the set of all paths, and Σ_{free} the set of all collision-free paths. Given two paths $\sigma_1, \sigma_2 \in \Sigma$, such that $\sigma_1(1) = \sigma_2(0)$, let $\sigma_1|\sigma_2 \in \Sigma$ denote their concatenation, i.e., $(\sigma_1|\sigma_2)(\tau) := \sigma_1(2\tau)$ for all $\tau \in [0, 1/2]$ and $(\sigma_1|\sigma_2)(\tau) := \sigma_2(2\tau - 1)$ for all $\tau \in (1/2, 1]$. Both Σ and Σ_{free} are closed under concatenation. Let $c : \Sigma \to \text{R}_{\geq 0}$ be a function, called the *cost function*, which assigns a strictly positive cost to all non-trivial collision-free paths (i.e., $c(\sigma) = 0$ if and only if $\sigma(\tau) = \sigma(0), \forall \tau \in [0,1]$). The cost function is assumed to be *monotonic*, in the sense that for all $\sigma_1, \sigma_2 \in \Sigma$, $c(\sigma_1) \leq c(\sigma_1|\sigma_2)$, and *bounded*, in the sense that there exists k_c such that $c(\sigma) \leq k_c \text{TV}(\sigma)$, $\forall \sigma \in \Sigma$.

The *optimality problem* of path planning asks for finding a feasible path with minimum cost:

Problem 3 (Optimal path planning) *Given a path planning problem* $(\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$ *and a cost function* $c : \Sigma \to \mathbb{R}_{\geq 0}$*, find a feasible path* σ^* *such that* $c(\sigma^*) = \min\{c(\sigma) : \sigma \text{ is feasible}\}$*. If no such path exists, report failure.*

2.2 Random Geometric Graphs

The objective of this section is to summarize some of the results on random geometric graphs that are available in the literature, and are relevant to the analysis of sampling-based path planning algorithms. In the remainder of this article, several connections are made between the theory of random geometric graphs and path-planning algorithms in robotics, providing insight on a number of issues, including, e.g., probabilistic completeness and asymptotic optimality, as well as technical tools to analyze the algorithms and establish their properties. In fact, it turns out that the data structures constructed by most sampling-based motion planning algorithms in the literature coincide, in the absence of obstacles, with standard models of random geometric graphs.

Random geometric graphs are in general defined as stochastic collections of points in a metric space, connected pairwise by edges if certain conditions (e.g., on the distance between the points) are satisfied. Such objects have been studied since their introduction by Gilbert (1961); see, e.g., Penrose (2003) and Balister et al. (2009a) for an overview of recent results. From the theoretical point of view, the study of random geometric graphs makes a connection between random graphs (Bollobás, 2001) and percolation theory (Bollobás and Riordan, 2006). On the application side, in recent years, random geometric graphs have attracted significant attention as models of *ad hoc* wireless networks (Gupta and Kumar, 1998, 2000).

Much of the literature on random geometric graphs deals with infinite graphs defined on unbounded domains, with vertices generated as a homogeneous Poisson point process. Recall that a Poisson random variable of parameter $\lambda \in \mathbb{R}_{>0}$ is an integer-valued random variable $\mathrm{Poisson}(\lambda) : \Omega \to \mathbb{N}_0$ such that $\mathbb{P}(\mathrm{Poisson}(\lambda) = k) = e^{-\lambda}\lambda^k/k!$. A homogeneous Poisson point process of intensity λ on $\mathbb{R}^d$ is a random countable set of points $\mathcal{P}^d_\lambda \subset \mathbb{R}^d$ such that, for any disjoint measurable sets $\mathcal{S}_1, \mathcal{S}_2 \subset \mathbb{R}^d$, $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$, the numbers of points of $\mathcal{P}^d_\lambda$ in each set are independent Poisson variables, i.e., $\mathrm{card}\left(\mathcal{P}^d_\lambda \cap \mathcal{S}_1\right) = \mathrm{Poisson}(\mu(\mathcal{S}_1)\lambda)$ and $\mathrm{card}\left(\mathcal{P}^d_\lambda \cap \mathcal{S}_2\right) = \mathrm{Poisson}(\mu(\mathcal{S}_2)\lambda)$. In particular, the intensity of a homogeneous Poisson point process can be interpreted as the expected number of points generated in the unit cube, i.e., $\mathbb{E}(\mathrm{card}\left(\mathcal{P}^d_\lambda \cap (0,1)^d\right)) = \mathbb{E}(\mathrm{Poisson}(\lambda)) = \lambda$.

Perhaps the most studied model of infinite random geometric graph is the following, introduced in Gilbert (1961), and often called Gilbert's disc model, or Boolean model:

Definition 4 (Infinite random r-disc graph) *Let* $\lambda, r \in \mathbb{R}_{>0}$*, and* $d \in \mathbb{N}$*. An infinite random* r*-disc graph* $G^{\mathrm{disc}}_\infty(\lambda, r)$ *in* d *dimensions is an infinite graph with vertices* $\{X_i\}_{i \in \mathbb{N}} = \mathcal{P}^d_\lambda$*, and such that* (X_i, X_j)*,* $i, j \in \mathbb{N}$*, is an edge if and only if* $\|X_i - X_j\| < r$*.*

A fundamental issue in infinite random graphs is whether the graph contains an infinite connected component, with non-zero probability. If it does, the random graph is said to *percolate*. Percolation is an important paradigm in statistical physics, with many applications in disparate fields such as material science, epidemiology, and microchip manufacturing, just to name a few (see, e.g., Sahimi, 1994).

Consider the infinite random r-disc graph, for $r = 1$, i.e., $G^{\mathrm{disc}}_\infty(\lambda, 1)$, and assume, without loss of generality, that the origin is one of the vertices of this graph. Let $p_k(\lambda)$ denote the probability that the connected component of $G^{\mathrm{disc}}_\infty(\lambda, 1)$ containing the origin contains k vertices, and define $p_\infty(\lambda)$ as $p_\infty(\lambda) = 1 - \sum_{k=1}^\infty p_k(\lambda)$. The function $p_\infty : \lambda \to p_\infty(\lambda)$ is monotone, and $p_\infty(0) = 0$ and $\lim_{\lambda \to \infty} p_\infty(\lambda) = 1$ (Penrose, 2003). A key result in percolation theory is that there exists a non-zero *critical intensity* λ_{c} defined as $\lambda_{\mathrm{c}} := \sup\{\lambda : p_\infty(\lambda) = 0\}$. In other words, for all $\lambda > \lambda_{\mathrm{c}}$, there is a non-zero probability that the origin is in an infinite connected component of $G_\infty^{\text{disc}}(\lambda, 1)$; moreover, under these conditions, the graph has precisely one infinite connected component, almost surely (Meester and Roy, 1996). The function p_∞ is continuous for all $\lambda \neq \lambda_\text{c}$: in other words, the graph undergoes a phase transition at the critical density λ_c, often also called the continuum percolation threshold (Penrose, 2003). The exact value of λ_c is not known; Meester and Roy provide $0.696 < \lambda_c < 3.372$ for $d = 2$ (Meester and Roy, 1996), and simulations suggest that $\lambda_c \approx 1.44$ (Quintanilla et al., 2000).

For many applications, including the ones in this article, models of finite graphs on a bounded domain are more relevant. Penrose introduced the following model (Penrose, 2003):

Definition 5 (Random r-disc graph) *Let $r \in \mathbb{R}_{>0}$, and $n, d \in \mathbb{N}$. A random r-disc graph $G^{\text{disc}}(n, r)$ in d dimensions is a graph whose n vertices, $\{X_1, X_2, \ldots, X_n\}$, are independent, uniformly distributed random variables in $(0, 1)^d$, and such that (X_i, X_j), $i, j \in \{1, \ldots, n\}$, $i \neq j$, is an edge if and only if $\|X_i - X_j\| < r$.*

For finite random geometric graph models, one is typically interested in whether a random geometric graph possesses certain properties asymptotically as n increases. Since the number of vertices is finite in random graphs, percolation can not be defined easily. In this case, percolation is studied in terms of the scaling of the number of vertices in the largest connected component with respect to the total number of vertices; in particular, a finite random geometric graph is said to percolate if it contains a "giant" connected component containing at least a constant fraction of all the nodes. As in the infinite case, percolation in finite random geometric graphs is often a phase transition phenomenon. In the case of random r-disc graphs,

Theorem 6 (Percolation of random r-disc graphs (Penrose, 2003)) *Let $G^{\text{disc}}(n, r)$ be a random r-disc graph in $d \geq 2$ dimensions, and let $N_{\max}(G^{\text{disc}}(n, r))$ be the number of vertices in its largest connected component. Then, almost surely,*

$$\lim_{n\to\infty} \frac{N_{\max}(G^{\text{disc}}(n, r_n))}{n} = 0, \qquad \text{if } r_n < (\lambda_\text{c}/n)^{1/d},$$

and

$$\lim_{n\to\infty} \frac{N_{\max}(G^{\text{disc}}(n, r))}{n} > 0, \qquad \text{if } r_n > (\lambda_\text{c}/n)^{1/d},$$

where λ_c is the continuum percolation threshold.

A random r-disc graph with $\lim_{n\to\infty} n r_n^d = \lambda \in (0, \infty)$ is said to operate in the *thermodynamic limit*. It is said to be in *subcritical* regime when $\lambda < \lambda_c$ and *supercritical* regime when $\lambda > \lambda_c$.

Another property of interest is connectivity. Clearly, connectivity implies percolation. Interestingly, emergence of connectivity in random geometric graphs is a phase transition phenomenon, as percolation. The following result is available in the literature:

Theorem 7 (Connectivity of random r-disc graphs (Penrose, 2003)) *Let $G^{\text{disc}}(n, r)$ be a random r-disc graph in d dimensions. Then,*

$$\lim_{n\to\infty} \mathbb{P}\left(\{G^{\text{disc}}(n, r) \text{ is connected }\}\right) = \begin{cases} 1, & \text{if } \zeta_d r^d > \log(n)/n, \\ 0, & \text{if } \zeta_d r^d < \log(n)/n, \end{cases}$$

where ζ_d is the volume of the unit ball in d dimensions.

Another model of random geometric graphs considers edges between k nearest neighbors. (Note that there are no ties, almost surely.) Both infinite and finite models are considered, as follows.

Definition 8 (Infinite random k-nearest neighbor graph) *Let $\lambda \in \mathbb{R}_{>0}$, and $d, k \in \mathbb{N}$. An infinite random k-nearest neighbor graph $G_{\infty}^{\mathrm{near}}(\lambda, k)$ in d dimensions is an infinite graph with vertices $\{X_i\}_{i\in\mathbb{N}} = \mathcal{P}_{\lambda}^{d}$, and such that (X_i, X_j), $i, j \in \mathbb{N}$, is an edge if X_j is among the k nearest neighbors of X_i, or if X_i is among the k nearest neighbors of X_j.*

Definition 9 (Random k-nearest neighbor graph) *Let $d, k, n \in \mathbb{N}$. A random k-nearest neighbor graph $G^{\mathrm{near}}(n, k)$ in d dimensions is a graph whose n vertices, $\{X_1, X_2, \ldots, X_n\}$, are independent, uniformly distributed random variables in $(0,1)^d$, and such that (X_i, X_j), $i, j \in \{1, \ldots, n\}$, $i \neq j$, is an edge if X_j is among the k nearest neighbors of X_i, or if X_i is among the k nearest neighbors of X_j.*

Percolation and connectivity for random k-nearest neighbor graphs exhibit phase transition phenomena, as in the random r-disc case. However, the results available in the literature are more limited. Results on percolation are only available for infinite graphs:

Theorem 10 (Percolation in infinite random k-nearest graphs (Balister et al., 2009a)) *Let $G_{\infty}^{\mathrm{near}}(\lambda, k)$ be an infinite random k-nearest neighbor graph in $d \geq 2$ dimensions. Then, there exists a constant $k_d^{\mathrm{p}} > 0$ such that*

$$\mathbb{P}\left(\{G_{\infty}^{\mathrm{near}}(1, k) \textit{ has an infinite component }\}\right) = \begin{cases} 1, & \textit{if } k \geq k_d^{\mathrm{p}}, \\ 0, & \textit{if } k < k_d^{\mathrm{p}}. \end{cases}$$

The value of k_d^{p} is not known. However, it is believed that $k_2^{\mathrm{p}} = 3$, and $k_d^{\mathrm{p}} = 2$ for all $d \geq 3$ (Balister et al., 2009a). It is known that percolation does not occur for $k = 1$ (Balister et al., 2009a).

Regarding connectivity of random k-nearest neighbor graphs, the only available results in the literature are not stated in terms of a given number of vertices: rather, the results are stated in terms of the restriction of a homogeneous Poisson point process to the unit cube. In other words, the vertices of the graph are obtained as $\{X_1, X_2, \ldots\} = \mathcal{P}_{\lambda}^{d} \cap (0,1)^d$. This is equivalent to setting the number of vertices as a Poisson random variable of parameter n, and then sampling the Poisson(n) vertices independently and uniformly in $(0,1)^d$:

Lemma 11 (Stoyan et al. (1995)) *Let $\{X_i\}_{i\in\mathbb{N}}$ be a sequence of points drawn independently and uniformly from $\mathcal{S} \subseteq \mathcal{X}$. Let* Poisson$(n)$ *be a Poisson random variable with parameter n. Then, $\{X_1, X_2, \ldots, X_{\mathrm{Poisson}(n)}\}$ is the restriction to $\mathcal{S}$ of a homogeneous Poisson point process with intensity $n/\mu(\mathcal{S})$.*

The main advantage in using such a model to generate the vertices of a random geometric graph is independence: in the Poisson case, the numbers of points in any two disjoint measurable regions $\mathcal{S}_1, \mathcal{S}_2 \subset [0,1]^d$, $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$, are independent Poisson random variables, with mean $\mu(\mathcal{S}_1)\lambda$ and $\mu(\mathcal{S}_2)\lambda$, respectively. These two random variables would not be independent if the total number of vertices were fixed a priori (also called a binomial point process). With some abuse of notation, such a random geometric graph model will be indicated as $G^{\mathrm{near}}(\mathrm{Poisson}(n), k)$.

Theorem 12 (Connectivity of random k-nearest graphs (Balister et al., 2009b; Xue and Kumar, 20 *Let $G^{\mathrm{near}}(\mathrm{Poisson}(n), k)$ indicate a k-nearest neighbor graph model in $d = 2$ dimensions, such that* *its vertices are generated using a Poisson point process of intensity* n*. Then, there exists a constant* $k_2^{\mathrm{c}} > 0$ *such that*

$$\lim_{n\to\infty} \mathbb{P}\left(\{G^{\mathrm{near}}(\mathrm{Poisson}(n), \lfloor k \log(n) \rfloor) \textit{ is connected } \}\right) = \begin{cases} 1, & \textit{if } k \geq k_2^{\mathrm{c}}, \\ 0, & \textit{if } k < k_2^{\mathrm{c}}. \end{cases}$$

The value of k_2^{c} is not known; the current best estimate is $0.3043 \leq k_2^{\mathrm{c}} \leq 0.5139$ (Balister et al., 2005).

Finally, the last model of random geometric graph that will be relevant for the analysis of the algorithms in this paper is the following:

Definition 13 (Online nearest neighbor graph) *Let* $d, n \in \mathbb{N}$*. An online nearest neighbor graph* $G^{\mathrm{ONN}}(n)$ *in* d *dimensions is a graph whose* n *vertices,* $(X_1, X_2, \ldots, X_n)$*, are independent, uniformly distributed random variables in* $(0,1)^d$*, and such that* (X_i, X_j)*,* $i, j \in \{1, \ldots, n\}$*,* $j > 1$*, is an edge if and only if* $\|X_i - X_j\| = \min_{1\leq k<j} \|X_k - X_j\|$*.*

Clearly, the online nearest neighbor graph is connected by construction, and trivially percolates. Recent results for this random geometric graph model include estimates of the total power-weighted edge length and an analysis of the vertex degree distribution, see, e.g., Wade (2009).

3 Algorithms

In this section, a number of sampling-based motion planning algorithms are introduced. First, some common primitive procedures are defined. Then, the PRM and the RRT algorithms are outlined, as they are representative of the major paradigms for sampling-based motion planning algorithms in the literature. Then, new algorithms, namely PRM* and RRT*, are introduced, as asymptotically optimal and computationally efficient versions of their "standard" counterparts.

3.1 Primitive Procedures

Before discussing the algorithms, it is convenient to introduce the primitive procedures that they rely on.

Sampling: Let `Sample` : $\omega \mapsto \{\texttt{Sample}_i(\omega)\}_{i\in\mathbb{N}_0} \subset \mathcal{X}$ be a map from Ω to sequences of points in $\mathcal{X}$, such that the random variables $\texttt{Sample}_i$, $i \in \mathbb{N}_0$, are independent and identically distributed (i.i.d.). For simplicity, the samples are assumed to be drawn from a uniform distribution, even though results extend naturally to any absolutely continuous distribution with density bounded away from zero on $\mathcal{X}$. It is convenient to consider another map, `SampleFree` : $\omega \mapsto \{\texttt{SampleFree}_i(\omega)\}_{i\in\mathbb{N}_0} \subset \mathcal{X}_{\mathrm{free}}$ that returns sequences of i.i.d. samples from $\mathcal{X}_{\mathrm{free}}$. For each $\omega \in \Omega$, the sequence $\{\texttt{SampleFree}_i(\omega)\}_{i\in\mathbb{N}_0}$ is the subsequence of $\{\texttt{Sample}_i(\omega)\}_{i\in\mathbb{N}_0}$ containing only the samples in $\mathcal{X}_{\mathrm{free}}$, i.e., $\{\texttt{SampleFree}_i(\omega)\}_{i\in\mathbb{N}_0} = \{\texttt{Sample}_i(\omega)\}_{i\in\mathbb{N}_0} \cap \mathcal{X}_{\mathrm{free}}$.

Nearest Neighbor: Given a graph $G = (V, E)$, where $V \subset \mathcal{X}$, a point $x \in \mathcal{X}$, the function `Nearest` : $(G, x) \mapsto v \in V$ returns the vertex in V that is "closest" to x in terms of a given distance function. In this paper, the Euclidean distance is used (see, e.g., LaValle and Kuffner (2001) for alternative choices), and hence

$$\texttt{Nearest}(G = (V, E), x) := \mathrm{argmin}_{v\in V} \|x - v\|.$$

A set-valued version of this function is also considered, `kNearest` : $(G, x, k) \mapsto \{v_1, v_2, \dots, v_k\}$, returning the k vertices in V that are nearest to x, according to the same distance function as above. (By convention, if the cardinality of V is less than k, then the function returns V.)

Near Vertices: Given a graph $G = (V, E)$, where $V \subset \mathcal{X}$, a point $x \in \mathcal{X}$, and a positive real number $r \in \mathbb{R}_{>0}$, the function `Near` : $(G, x, r) \mapsto V' \subseteq V$ returns the vertices in V that are contained in a ball of radius r centered at x, i.e.,

$$\texttt{Near}(G = (V, E), x, r) := \{v \in V : v \in \mathcal{B}_{x,r}\}.$$

Steering: Given two points $x, y \in \mathcal{X}$, the function `Steer` : $(x, y) \mapsto z$ returns a point $z \in \mathcal{X}$ such that z is "closer" to y than x is. Throughout the paper, the point z returned by the function `Steer` will be such that z minimizes $\|z - y\|$ while at the same time maintaining $\|z - x\| \leq \eta$, for a prespecified $\eta > 0$,[1] i.e.,

$$\texttt{Steer}(x, y) := \mathrm{argmin}_{z \in \mathcal{B}_{x,\eta}} \|z - y\|.$$

Collision Test: Given two points $x, x' \in \mathcal{X}$, the Boolean function `CollisionFree`(x, x') returns `True` if the line segment between x and x' lies in $\mathcal{X}_{\mathrm{free}}$, i.e., $[x, x'] \subset \mathcal{X}_{\mathrm{free}}$, and `False` otherwise.

3.2 Existing Algorithms

Next, some of the sampling-based algorithms available in the literature are outlined. For convenience, inputs and outputs of the algorithms are not shown explicitly, but are as follows. All algorithms take as input a path planning problem $(\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$, an integer $n \in \mathbb{N}$, and a cost function $c : \Sigma \to \mathbb{R}_{\geq 0}$, if appropriate. These inputs are shared with functions and procedures called within the algorithms. All algorithms return a graph $G = (V, E)$, where $V \subset \mathcal{X}_{\mathrm{free}}$, $\mathrm{card}\,(V) \leq n+1$, and $E \in \mathrm{V} \times \mathrm{V}$. The solution of the path planning problem can be easily computed from such a graph, e.g., using standard shortest-path algorithms.

Probabilistic RoadMaps (PRM): The Probabilistic RoadMaps algorithm is primarily aimed at multi-query applications. In its basic version, it consists of a pre-processing phase, in which a roadmap is constructed by attempting connections among n randomly-sampled points in $\mathcal{X}_{\mathrm{free}}$, and a query phase, in which paths connecting initial and final conditions through the roadmap are sought. "Expansion" heuristics for enhancing the roadmap's connectivity are available in the literature (Kavraki et al., 1996) but have no impact on the analysis in this paper, and will not be discussed.

The pre-processing phase, outlined in Algorithm 1, begins with an empty graph. At each iteration, a point $x_{\mathrm{rand}} \in \mathcal{X}_{\mathrm{free}}$ is sampled, and added to the vertex set V. Then, connections are attempted between x_{rand} and other vertices in V within a ball of radius r centered at x_{rand}, in order of increasing distance from x_{rand}, using a simple local planner (e.g., straight-line connection). Successful (i.e., collision-free) connections result in the addition of a new edge to the edge set E. To avoid unnecessary computations (since the focus of the algorithm is establishing connectivity), connections between x_{rand} and vertices in the same connected component are avoided. Hence, the roadmap constructed by PRM is a forest, i.e., a collection of trees.

---

[1] This steering procedure is used widely in the robotics literature, since its introduction in Kuffner and LaValle (2000). Our results also extend to the Rapidly-exploring Random Dense Trees (see, e.g., LaValle, 2006), which are slightly modified versions of the RRTs that do not require tuning any prespecified parameters such as η in this case.

Algorithm 1: PRM (preprocessing phase)

```
1 $V \leftarrow \emptyset$; $E \leftarrow \emptyset$;
2 for $i = 0, \ldots, n$ do
3     $x_{\text{rand}} \leftarrow$ SampleFree$_i$;
4     $U \leftarrow$ Near$(G = (V, E), x_{\text{rand}}, r)$ ;
5     $V \leftarrow V \cup \{x_{\text{rand}}\}$;
6     foreach $u \in U$, in order of increasing $\|u - x_{\text{rand}}\|$, do
7         if $x_{\text{rand}}$ and $u$ are not in the same connected component of $G = (V, E)$ then
8             if CollisionFree$(x_{\text{rand}}, u)$ then $E \leftarrow E \cup \{(x_{\text{rand}}, u), (u, x_{\text{rand}})\}$;
9 return $G = (V, E)$;
```

Analysis results in the literature are only available for a "simplified" version of the PRM algorithm (Kavraki et al., 1998), referred to as sPRM in this paper. The simplified algorithm initializes the vertex set with the initial condition, samples n points from $\mathcal{X}_{\text{free}}$, and then attempts to connect points within a distance r, i.e., using a similar logic as PRM, with the difference that connections between vertices in the same connected component are allowed. Notice that in the absence of obstacles, i.e., if $\mathcal{X}_{\text{free}} = \mathcal{X}$, the roadmap constructed in this way is a random r-disc graph.

Algorithm 2: sPRM

```
1 $V \leftarrow \{x_{\text{init}}\} \cup \{$SampleFree$_i\}_{i=1,\ldots,n}$; $E \leftarrow \emptyset$;
2 foreach $v \in V$ do
3     $U \leftarrow$ Near$(G = (V, E), v, r) \setminus \{v\}$;
4     foreach $u \in U$ do
5         if CollisionFree$(v, u)$ then $E \leftarrow E \cup \{(v, u), (u, v)\}$
6 return $G = (V, E)$;
```

Practical implementation of the (s)PRM algorithm have often considered different choices for the set U of vertices to which connections are attempted (i.e., line 4 in Algorithm 1, and line 3 in Algorithm 2). In particular, the following criteria are of particular interest:

- ***k*-Nearest (s)PRM**: Choose the nearest k neighbors to the vertex under consideration, for a given k (a typical value is reported as $k = 15$ (LaValle, 2006)). In other words, $U \leftarrow$ `kNearest`$(G = (V, E), x_{\text{rand}}, k)$ in line 4 of Algorithm 1 and $U \leftarrow$ `kNearest`$(G = (V, E), v, k) \setminus \{v\}$ in line 3 of Algorithm 2. The roadmap constructed in this way in an obstacle-free environment is a random k-nearest graph.

- **Bounded-degree (s)PRM**: For any fixed r, the average number of connections attempted at each iteration is proportional to the number of vertices in V, and can result in an excessive computational burden for large n. To address this, an upper bound k can be imposed on the cardinality of the set U (a typical value is reported as $k = 20$ (LaValle, 2006)). In other words, $U \leftarrow$ `Near`$(G, x_{\text{rand}}, r) \cap$ `kNearest`(G, x_{rand}, k) in line 4 of Algorithm 1, and $U \leftarrow ($`Near`$(G, v, r) \cap$ `kNearest`$(G, v, k)) \setminus \{v\}$ in line 3 of Algorithm 2.

- **Variable-radius (s)PRM**: Another option to maintain the degree of the vertices in the roadmap small is to make the connection radius r a function of n, as opposed to a fixed parameter. However, there are no clear indications in the literature on the appropriate functional relationship between r and n.

Rapidly-exploring Random Trees (RRT): The Rapidly-exploring Random Tree algorithm is primarily aimed at single-query applications. In its basic version, the algorithm incrementally builds a tree of feasible trajectories, rooted at the initial condition. An outline of the algorithm is given in Algorithm 3. The algorithm is initialized with a graph that includes the initial state as its single vertex, and no edges. At each iteration, a point $x_{\mathrm{rand}} \in \mathcal{X}_{\mathrm{free}}$ is sampled. An attempt is made to connect the nearest vertex $v \in V$ in the tree to the new sample. If such a connection is successful, x_{rand} is added to the vertex set, and (v, x_{rand}) is added to the edge set. In the original version of this algorithm, the iteration is stopped as soon as the tree contains a node in the goal region. In this paper, for consistency with the other algorithms (e.g., PRM), the iteration is performed n times. In the absence of obstacles, i.e., if $\mathcal{X}_{\mathrm{free}} = \mathcal{X}$, the tree constructed in this way is an online nearest neighbor graph.

```
Algorithm 3: RRT
$V \leftarrow \{x_{\mathrm{init}}\}$; $E \leftarrow \emptyset$;
for $i = 1, \ldots, n$ do
    $x_{\mathrm{rand}} \leftarrow$ SampleFree$_i$;
    $x_{\mathrm{nearest}} \leftarrow$ Nearest$(G = (V, E), x_{\mathrm{rand}})$;
    $x_{\mathrm{new}} \leftarrow$ Steer$(x_{\mathrm{nearest}}, x_{\mathrm{rand}})$ ;
    if ObtacleFree$(x_{\mathrm{nearest}}, x_{\mathrm{new}})$ then
        $V \leftarrow V \cup \{x_{\mathrm{new}}\}$; $E \leftarrow E \cup \{(x_{\mathrm{nearest}}, x_{\mathrm{new}})\}$ ;
return $G = (V, E)$;
```

A variant of RRT consists of growing two trees, respectively rooted at the initial state, and at a state in the goal set. To highlight the fact that the sampling procedure must not necessarily be stochastic, the algorithm is also referred to as Rapidly-exploring Dense Trees (RDT) (LaValle, 2006).

3.3 Proposed algorithms

In this section, the new algorithms considered in this paper are presented. These algorithms are proposed as asymptotically optimal and computationally efficient versions of their "standard" counterparts, as will be made clear through the analysis in the next section. Input and output data are the same as in the algorithms introduced in Section 3.2.

Optimal Probabilistic RoadMaps (PRM*): In the standard PRM algorithm, as well as in its simplified "batch" version considered in this paper, connections are attempted between roadmap vertices that are within a fixed radius r from one another. The constant r is thus a parameter of PRM. The proposed algorithm—shown in Algorithm 4—is similar to sPRM, with the only difference being that the connection radius r is chosen as a function of n, i.e., $r = r(n) := \gamma_{\mathrm{PRM}}(\log(n)/n)^{1/d}$, where $\gamma_{\mathrm{PRM}} > \gamma^*_{\mathrm{PRM}} = 2(1 + 1/d)^{1/d}\,(\mu(\mathcal{X}_{\mathrm{free}})/\zeta_d)^{1/d}$, d is the dimension of the space $\mathcal{X}$, $\mu(\mathcal{X}_{\mathrm{free}})$ denotes the Lebesgue measure (i.e., volume) of the obstacle-free space, and ζ_d is the volume of the unit ball in the d-dimensional Euclidean space. Clearly, the connection radius decreases with the number of samples. The rate of decay is such that the average number of connections attempted from a roadmap vertex is proportional to $\log(n)$.

Note that in the discussion of variable-radius PRM in LaValle (2006), it is suggested that the radius be chosen as a function of sample dispersion. (Recall that the dispersion of a point set contained in a bounded set $\mathcal{S} \subset \mathbb{R}^d$ is the radius of the largest empty ball centered in $\mathcal{S}$.) Indeed, the dispersion of a set of n random points sampled uniformly and independently in a bounded set is $O((\log(n)/n)^{1/d})$ (Niederreiter, 1992), which is precisely the rate at which the connection radius is scaled in the PRM* algorithm.

Algorithm 4: PRM*

```
1 $V \leftarrow \{x_{\mathrm{init}}\} \cup \{\mathtt{SampleFree}_i\}_{i=1,\ldots,n}$; $E \leftarrow \emptyset$;
2 foreach $v \in V$ do
3     $U \leftarrow \mathtt{Near}(G = (V, E), v, \gamma_{\mathrm{PRM}}(\log(n)/n)^{1/d}) \setminus \{v\}$;
4     foreach $u \in U$ do
5         if $\mathtt{CollisionFree}(v, u)$ then $E \leftarrow E \cup \{(v, u), (u, v)\}$
6 return $G = (V, E)$;
```

Another version of the algorithm, called k-nearest PRM*, can be considered, motivated by the k-nearest PRM implementation previously mentioned, whereby the number k of nearest neighbors to be considered is not a constant, but is chosen as a function of the cardinality of the roadmap n. More precisely, $k(n) := k_{\mathrm{PRM}} \log(n)$, where $k_{\mathrm{PRM}} > k^*_{\mathrm{PRM}} = e\,(1 + 1/d)$, and $U \leftarrow \mathtt{kNearest}(G = (V, E), v, k_{\mathrm{PRM}} \log(n)) \setminus \{v\}$ in line 3 of Algorithm 4.

Note that k^*_{PRM} is a constant that only depends on d, and does not otherwise depend on the problem instance, unlike γ^*_{PRM}. Moreover, $k_{\mathrm{PRM}} = 2e$ is a valid choice for all problem instances.

Rapidly-exploring Random Graph (RRG): The Rapidly-exploring Random Graph algorithm was introduced as an incremental (as opposed to batch) algorithm to build a *connected* roadmap, possibly containing cycles. The RRG algorithm is similar to RRT in that it first attempts to connect the nearest node to the new sample. If the connection attempt is successful, the new node is added to the vertex set. However, RRG has the following difference. Every time a new point x_{new} is added to the vertex set V, then connections are attempted from all other vertices in V that are within a ball of radius $r(\mathrm{card}\,(V)) = \min\{\gamma_{\mathrm{RRG}}(\log(\mathrm{card}\,(V))/\,\mathrm{card}\,(V))^{1/d}, \eta\}$, where η is the constant appearing in the definition of the local steering function, and $\gamma_{\mathrm{RRG}} > \gamma^*_{\mathrm{RRG}} = 2\,(1 + 1/d)^{1/d}\,(\mu(\mathcal{X}_{\mathrm{free}})/\zeta_d)^{1/d}$. For each successful connection, a new edge is added to the edge set E. Hence, it is clear that, for the same sampling sequence, the RRT graph (a directed tree) is a subgraph of the RRG graph (an undirected graph, possibly containing cycles). In particular, the two graphs share the same vertex set, and the edge set of the RRT graph is a subset of that of the RRG graph.

Another version of the algorithm, called k-nearest RRG, can be considered, in which connections are sought to k nearest neighbors, with $k = k(\mathrm{card}\,(V)) := k_{\mathrm{RRG}} \log(\mathrm{card}\,(V))$, where $k_{\mathrm{RRG}} > k^*_{\mathrm{RRG}} = e\,(1 + 1/d)$, and $X_{\mathrm{near}} \leftarrow \mathtt{kNearest}(G = (V, E), x_{\mathrm{new}}, k_{\mathrm{RRG}} \log(\mathrm{card}\,(V)))$, in line 7 of Algorithm 5.

Note that k^*_{RRG} is a constant that depends only on d, and does not depend otherwise on the problem instance, unlike γ^*_{RRG}. Moreover, $k_{\mathrm{RRG}} = 2e$ is a valid choice for all problem instances.

Optimal RRT (RRT*): Maintaining a tree structure rather than a graph is not only economical in terms of memory requirements, but may also be advantageous in some applications, due to, for instance, relatively easy extensions to motion planning problems with differential constraints, or

Algorithm 5: RRG

```
$V \leftarrow \{x_{\mathrm{init}}\}$; $E \leftarrow \emptyset$;
for $i = 1, \ldots, n$ do
    $x_{\mathrm{rand}} \leftarrow$ SampleFree$_i$;
    $x_{\mathrm{nearest}} \leftarrow$ Nearest$(G = (V, E), x_{\mathrm{rand}})$;
    $x_{\mathrm{new}} \leftarrow$ Steer$(x_{\mathrm{nearest}}, x_{\mathrm{rand}})$ ;
    if ObtacleFree$(x_{\mathrm{nearest}}, x_{\mathrm{new}})$ then
        $X_{\mathrm{near}} \leftarrow$ Near$(G = (V, E), x_{\mathrm{new}}, \min\{\gamma_{\mathrm{RRG}}(\log(\mathrm{card}\,(V))/\,\mathrm{card}\,(V))^{1/d}, \eta\})$ ;
        $V \leftarrow V \cup \{x_{\mathrm{new}}\}$; $E \leftarrow E \cup \{(x_{\mathrm{nearest}}, x_{\mathrm{new}}), (x_{\mathrm{new}}, x_{\mathrm{nearest}})\}$ ;
        foreach $x_{\mathrm{near}} \in X_{\mathrm{near}}$ do
            if CollisionFree$(x_{\mathrm{near}}, x_{\mathrm{new}})$ then $E \leftarrow E \cup \{(x_{\mathrm{near}}, x_{\mathrm{new}}), (x_{\mathrm{new}}, x_{\mathrm{near}})\}$
return $G = (V, E)$;
```

to cope with modeling errors. The RRT* algorithm is obtained by modifying RRG in such a way that formation of cycles is avoided, by removing "redundant" edges, i.e., edges that are not part of a shortest path from the root of the tree (i.e., the initial state) to a vertex. Since the RRT and RRT* graphs are directed trees with the same root and vertex set, and edge sets that are subsets of that of RRG, this amounts to a "rewiring" of the RRT tree, ensuring that vertices are reached through a minimum-cost path.

Before discussing the algorithm, it is necessary to introduce a few new functions. Given two points $x_1, x_2 \in \mathbb{R}^d$, let `Line`$(x_1, x_2) : [0, s] \to \mathcal{X}$ denote the straight-line path from x_1 to x_2. Given a tree $G = (V, E)$, let `Parent` $: V \to V$ be a function that maps a vertex $v \in V$ to the unique vertex $u \in V$ such that $(u, v) \in E$. By convention, if $v_0 \in V$ is the root vertex of G, `Parent`$(v_0) = v_0$. Finally, let `Cost` $: V \to \mathbb{R}_{\geq 0}$ be a function that maps a vertex $v \in V$ to the cost of the unique path from the root of the tree to v. For simplicity, in stating the algorithm we will assume an additive cost function, so that `Cost`$(v) =$ `Cost`$($`Parent`$(v)) + c($`Line`$($`Parent`$(v), v))$, although this is not necessary for the analysis in the next section. By convention, if $v_0 \in V$ is the root vertex of G, then `Cost`$(v_0) = 0$.

The RRT* algorithm, shown in Algorithm 6, adds points to the vertex set V in the same way as RRT and RRG. It also considers connections from the new vertex x_{new} to vertices in X_{near}, i.e., other vertices that are within distance $r(\mathrm{card}\,(V)) = \min\{\gamma_{\mathrm{RRT}^*}(\log(\mathrm{card}\,(V))/\,\mathrm{card}\,(V))^{1/d}, \eta\}$ from x_{new}. However, not all feasible connections result in new edges being inserted in the edge set E. In particular, (i) an edge is created from the vertex in X_{near} that can be connected to x_{new} along a path with minimum cost, and (ii) new edges are created from x_{new} to vertices in X_{near}, if the path through x_{new} has lower cost than the path through the current parent; in this case, the edge linking the vertex to its current parent is deleted, to maintain the tree structure.

Another version of the algorithm, called k-nearest RRT*, can be considered, in which connections are sought to k nearest neighbors, with $k(\mathrm{card}\,(V)) = k_{\mathrm{RRG}} \log(\mathrm{card}\,(V))$, and $X_{\mathrm{near}} \leftarrow$ `kNearest`$(G = (V, E), x_{\mathrm{new}}, k_{\mathrm{RRG}} \log(i))$, in line 7 of Algorithm 6.

4 Analysis

In this section, a number of results concerning the probabilistic completeness, asymptotic optimality, and complexity of the algorithms in Section 3 are presented.

The return value of Algorithms 1-6 is a graph. Since the sampling procedure `SampleFree` is

Algorithm 6: RRT*

```
$V \leftarrow \{x_{\mathrm{init}}\}$; $E \leftarrow \emptyset$;
for $i = 1, \ldots, n$ do
    $x_{\mathrm{rand}} \leftarrow$ SampleFree$_i$;
    $x_{\mathrm{nearest}} \leftarrow$ Nearest$(G = (V, E), x_{\mathrm{rand}})$;
    $x_{\mathrm{new}} \leftarrow$ Steer$(x_{\mathrm{nearest}}, x_{\mathrm{rand}})$ ;
    if ObtacleFree$(x_{\mathrm{nearest}}, x_{\mathrm{new}})$ then
        $X_{\mathrm{near}} \leftarrow$ Near$(G = (V, E), x_{\mathrm{new}}, \min\{\gamma_{\mathrm{RRT}^*}(\log(\mathrm{card}\,(V))/\,\mathrm{card}\,(V))^{1/d}, \eta\})$ ;
        $V \leftarrow V \cup \{x_{\mathrm{new}}\}$;
        $x_{\mathrm{min}} \leftarrow x_{\mathrm{nearest}}$; $c_{\mathrm{min}} \leftarrow$ Cost$(x_{\mathrm{nearest}}) + c($Line$(x_{\mathrm{nearest}}, x_{\mathrm{new}}))$;
        foreach $x_{\mathrm{near}} \in X_{\mathrm{near}}$ do                // Connect along a minimum-cost path
            if CollisionFree$(x_{\mathrm{near}}, x_{\mathrm{new}}) \wedge$ Cost$(x_{\mathrm{near}}) + c($Line$(x_{\mathrm{near}}, x_{\mathrm{new}})) < c_{\mathrm{min}}$ then
                $x_{\mathrm{min}} \leftarrow x_{\mathrm{near}}$; $c_{\mathrm{min}} \leftarrow$ Cost$(x_{\mathrm{near}}) + c($Line$(x_{\mathrm{near}}, x_{\mathrm{new}}))$
        $E \leftarrow E \cup \{(x_{\mathrm{min}}, x_{\mathrm{new}})\}$;
        foreach $x_{\mathrm{near}} \in X_{\mathrm{near}}$ do                // Rewire the tree
            if CollisionFree$(x_{\mathrm{new}}, x_{\mathrm{near}}) \wedge$ Cost$(x_{\mathrm{new}}) + c($Line$(x_{\mathrm{new}}, x_{\mathrm{near}})) <$ Cost$(x_{\mathrm{near}})$
            then  $x_{\mathrm{parent}} \leftarrow$ Parent$(x_{\mathrm{near}})$;
                $E \leftarrow (E \setminus \{(x_{\mathrm{parent}}, x_{\mathrm{near}})\}) \cup \{(x_{\mathrm{new}}, x_{\mathrm{near}})\}$
return $G = (V, E)$;
```

stochastic, the returned graph is in fact a random variable.[2] Since the sampling procedure is modeled as a map from the sample space Ω to infinite sequences in $\mathcal{X}$, sets of vertices and edges of the graphs maintained by the algorithms can be defined as functions from the sample space Ω to appropriate sets. More precisely, let ALG be a label indicating one of the algorithms in Section 3, and let $\{V_i^{\mathrm{ALG}}(\omega)\}_{i \in \mathbb{N}}$ and $\{E_i^{\mathrm{ALG}}(\omega)\}_{i \in \mathbb{N}}$ be, respectively, the sets of vertices and edges in the graph returned by algorithm ALG, indexed by the number of samples, for a particular realization of the sample sequence. (In other words, these are sequences of functions defined from Ω into finite subsets of $\mathcal{X}_{\mathrm{free}}$ or $\mathcal{X}_{\mathrm{free}} \times \mathcal{X}_{\mathrm{free}}$.) Similarly, let $G_i^{\mathrm{ALG}} = (V_i^{\mathrm{ALG}}, E_i^{\mathrm{ALG}})$. (The label ALG will be at times omitted when the algorithm being used is clear from the context.)

All algorithms considered in the paper are sound, in the sense that they only return graphs with vertices and edges representing points and paths in $\mathcal{X}_{\mathrm{free}}$.This statement can be easily verified by inspection of the algorithms in Section 3.

4.1 Probabilistic Completeness

In this section, the feasibility problem is considered, and the (probabilistic) completeness properties of the algorithms in Section 3 are analyzed. First, some preliminary definitions are given, followed by a definition of probabilistic completeness. Then, completeness properties of various sampling-based motion planning algorithms are stated.

Let $\delta > 0$ be a real number. A state $x \in \mathcal{X}_{\mathrm{free}}$ is said to be a δ-*interior state* of $\mathcal{X}_{\mathrm{free}}$, if the closed ball of radius δ centered at x lies entirely inside $\mathcal{X}_{\mathrm{free}}$. The δ-*interior* of $\mathcal{X}_{\mathrm{free}}$, denoted as $\mathrm{int}_\delta(\mathcal{X}_{\mathrm{free}})$, is defined as the collection of all δ-interior states, i.e., $\mathrm{int}_\delta(\mathcal{X}_{\mathrm{free}}) := \{x \in \mathcal{X}_{\mathrm{free}} \,|\, \mathcal{B}_{x,\delta} \subseteq \mathcal{X}_{\mathrm{free}}\}$. In

---

[2] We will not address the case in which the sampling procedure is deterministic, but refer the reader to LaValle et al. (2004), which contains an in-depth discussion of the relative merits of randomness and determinism in sampling-based motion planning algorithms.

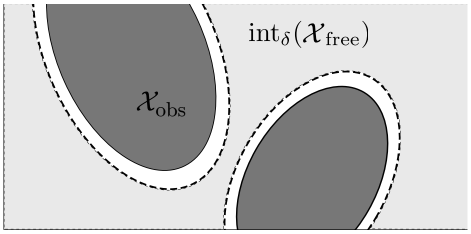
$\mathrm{int}_\delta(\mathcal{X}_{\mathrm{free}})$
$\mathcal{X}_{\mathrm{obs}}$

Figure 1: An illustration of the δ-interior of $\mathcal{X}_{\mathrm{free}}$. The obstacle region $\mathcal{X}_{\mathrm{obs}}$ is shown in dark grey and the δ-interior of $\mathcal{X}_{\mathrm{free}}$ is shown in light grey. The distance between the dashed boundary of $\mathrm{int}_\delta(\mathcal{X}_{\mathrm{free}})$ and the solid boundary of $\mathcal{X}_{\mathrm{free}}$ is precisely δ.

other words, the δ-interior of $\mathcal{X}_{\mathrm{free}}$ is the set of all states that are at least a distance δ away from any point in the obstacle set (see Figure 1). A collision-free path $\sigma : [0,1] \to \mathcal{X}_{\mathrm{free}}$ is said to have *strong δ-clearance*, if σ lies entirely inside the δ-interior of $\mathcal{X}_{\mathrm{free}}$, i.e., $\sigma(\tau) \in \mathrm{int}_\delta(\mathcal{X}_{\mathrm{free}})$ for all $\tau \in [0,1]$. A path planning problem $(\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$ is said to be *robustly feasible* if there exists a path with strong δ-clearance, for some $\delta > 0$, that solves it. In terms of the notation used in this paper, the notion of probabilistic completeness can be stated as follows.

Definition 14 (Probabilistic Completeness) *An algorithm ALG is probabilistically complete, if, for any robustly feasible path planning problem $(\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$,*

$$\liminf_{n\to\infty} \mathbb{P}\left(\{\exists x_{\mathrm{goal}} \in V_n^{\mathrm{ALG}} \cap \mathcal{X}_{\mathrm{goal}} \textit{ such that } x_{\mathrm{init}} \textit{ is connected to } x_{\mathrm{goal}} \textit{ in } G_n^{\mathrm{ALG}}\ \}\right) = 1.$$

If an algorithm is probabilistically complete, and the path planning problem is robustly feasible, the limit

$$\lim\nolimits_{n\to\infty} \mathbb{P}\left(\{\exists x_{\mathrm{goal}} \in V_n^{\mathrm{ALG}} \cap \mathcal{X}_{\mathrm{goal}} \text{ such that } x_{\mathrm{init}} \text{ is connected to } x_{\mathrm{goal}} \text{ in } G_n^{\mathrm{ALG}}\ \}\right)$$

exists and is equal to 1. On the other hand, the same limit is equal to zero for any sampling-based algorithm (including probabilistically complete ones) if the problem is not robustly feasible, unless the samples are drawn from a singular distribution adapted to the problem.

It is known from the literature that the sPRM and RRT algorithms are probabilistically complete, and that the probability of finding a solution if one exists approaches one exponentially fast with the number of vertices in the graph returned by the algorithms. In other words,

Theorem 15 (Probabilistic completeness of sPRM (Kavraki et al., 1998)) *Consider a robustly feasible path planning problem $(\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$. There exist constants $a > 0$ and $n_0 \in \mathbb{N}$, dependent only on $\mathcal{X}_{\mathrm{free}}$ and $\mathcal{X}_{\mathrm{goal}}$, such that*

$$\mathbb{P}\left(\left\{\exists\, x_{\mathrm{goal}} \in V_n^{\mathrm{sPRM}} \cap \mathcal{X}_{\mathrm{goal}} : x_{\mathrm{goal}} \textit{ is connected to } x_{\mathrm{init}} \textit{ in } G_n^{\mathrm{sPRM}}\right\}\right) > 1 - e^{-a\,n}, \quad \forall n > n_0.$$

Theorem 16 (Probabilistic Completeness of RRT (LaValle and Kuffner, 2001)) *Consider a robustly feasible path planning problem $(\mathcal{X}_{\mathrm{free}},\ x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$. There exist constants $a > 0$ and $n_0 \in \mathbb{N}$, both dependent only on $\mathcal{X}_{\mathrm{free}}$ and $\mathcal{X}_{\mathrm{goal}}$, such that*

$$\mathbb{P}\left(\left\{V_n^{\mathrm{RRT}} \cap \mathcal{X}_{\mathrm{goal}} \neq \emptyset\right\}\right) > 1 - e^{-a\,n}, \quad \forall n > n_0.$$

On the other hand, the probabilistic completeness results do not necessarily extend to the heuristics used in practical implementations of the (s)PRM algorithm, as detailed in Section 3. For example, consider the k-nearest sPRM algorithm, where $k = 1$. That is, each vertex is connected to its nearest neighbor and the resulting undirected graph is returned as the output. This sPRM algorithm will be called the 1-nearest sPRM, and indicated with the label 1PRM. The RRT algorithm can be thought of as the incremental version of the 1-nearest sPRM algorithm: the RRT algorithm also connects each sample to its nearest neighbor, but forces connectivity of the graph by an incremental construction. The following theorem shows that the 1-nearest sPRM algorithm is not probabilistically complete, although the RRT is (see Theorem 16). Furthermore, the probability that it *fails* to find a path converges to one as the number of samples approaches infinity.

Theorem 17 (Incompleteness of k-nearest sPRM for $k = 1$) *The k-nearest sPRM algorithm is not probabilistically complete for $k = 1$. Furthermore,*

$$\lim_{n\to\infty} \mathbb{P}\left(\{\exists x_{\mathrm{goal}} \in V_n^{\mathrm{1PRM}} \cap \mathcal{X}_{\mathrm{goal}} \text{ such that } x_{\mathrm{init}} \text{ is connected to } x_{\mathrm{goal}} \text{ in } G_n^{\mathrm{ALG}}\ \}\right) = 0.$$

The proof of this theorem requires two intermediate results that are provided below. For simplicity of presentation, consider the case when $\mathcal{X}_{\mathrm{free}} = \mathcal{X}$. Let $G_n^{\mathrm{1PRM}} = (V_n^{\mathrm{1PRM}}, E_n^{\mathrm{1PRM}})$ denote the graph returned by the 1-nearest sPRM algorithm, when the algorithm is run with n samples. Let L_n denote the total length of all the edges present in G_n^{1PRM}. Recall that ζ_d denotes the volume of the unit ball in the d-dimensional Euclidean space. Let ζ_d' denote the volume of the union of two unit balls whose centers are a unit distance apart.

Lemma 18 (Total length of the 1-nearest neighbor graph (Wade, 2007)) *For all $d \geq 2$, $L_n/n^{1-1/d}$ converges to a constant in mean square, i.e.,*

$$\lim_{n\to\infty} \mathbb{E}\left[\left(\frac{L_n}{n^{1-1/d}} - \left(1 + \frac{1}{d}\right)\left(\frac{1}{\zeta_d} - \frac{\zeta_d}{2\,(\zeta_d')^{1+1/d}}\right)\right)^2\right] = 0.$$

Proof This lemma is a direct consequence of Theorem 3 of Wade (2007). ☐

Let N_n denote the number of connected components of G_n^{1PRM}.

Lemma 19 (Number of connected components of the 1-nearest neighbor graph) *For all $d \geq 2$, N_n/n converges to a constant in mean square, i.e.,*

$$\lim_{n\to\infty} \mathbb{E}\left[\left(\frac{N_n}{n} - \frac{\zeta_d}{2\,\zeta_d'}\right)^2\right] = 0.$$

Proof A reciprocal pair is a pair of vertices each of which is the other one's nearest neighbor. In a graph formed by connecting each vertex to its nearest neighbor, any connected component includes exactly one reciprocal pair whenever the number of vertices is greater than 2 (see, e.g., Eppstein et al., 1997). The number of reciprocal pairs in such a graph was shown to converge to $\zeta_d/(2\zeta_d')$ in mean square in Henze (1987) (see also Remark 2 in Wade (2007)). ☐

Proof of Theorem 17 Let $\widetilde{L}_n$ denote the average length of a connected component in G_n^{1PRM}, i.e., $\widetilde{L}_n = L_n/N_n$. Let L_n' denote the length of the connected component that includes x_{init}. Since the samples are drawn independently and uniformly, the random variables $\widetilde{L}_n$ and L_n' have the same distribution (although they are clearly dependent). Let γ_L denote the constant that $L_n/n^{1-1/d}$ converges to (see Lemma 18). Similarly, let γ_N denote the constant that N_n/n converges to (see Lemma 19).

Recall that convergence in mean square implies convergence in probability and hence convergence in distribution (Grimmett and Stirzaker, 2001). Since both $L_n/n^{1-1/d}$ and N_n/n converge in mean square to constants and $\mathbb{P}(\{N_n = 0\}) = 0$ for all $n \in \mathbb{N}$, by Slutsky's theorem (Resnick, 1999), $n^{1/d}\,\widetilde{L}_n = \frac{L_n/n^{1-1/d}}{N_n/n}$ converges to $\gamma := \gamma_L/\gamma_N$ in distribution. In this case, it also converges in probability, since γ is a constant (Grimmett and Stirzaker, 2001). Then, $n^{1/d}\,L'_n$ also converges to γ in probability, since $\widetilde{L}_n$ and L'_n are identically distributed for all $n \in \mathbb{N}$. Thus, L'_n converges to 0 in probability, i.e., $\lim_{n\to\infty} \mathbb{P}\left(\{L'_n > \epsilon\}\right) = 0$, for all $\epsilon > 0$.

Let $\epsilon > 0$ be such that $\epsilon < \inf_{x \in \mathcal{X}_{\text{goal}}} \|x - x_{\text{init}}\|$. Let A_n denote the event that the graph returned by the 1-nearest sPRM algorithm contains a feasible path, i.e., one that starts from x_{init} and reaches the goal region Clearly, the event $\{L'_n > \epsilon\}$ occurs whenever A_n does, i.e., $A_n \subseteq \{L'_n > \epsilon\}$. Then, $\mathbb{P}(A_n) \leq \mathbb{P}(\{L'_n > \epsilon\})$. Taking the limit superior of both sides

$$\liminf_{n\to\infty} \mathbb{P}(A_n) \;\leq\; \limsup_{n\to\infty} \mathbb{P}(A_n) \;\leq\; \limsup_{n\to\infty} \mathbb{P}(\{L'_n > \epsilon\}) \;=\; 0.$$

In other words, the limit $\lim_{n\to\infty} \mathbb{P}(A_n)$ exists and is equal zero. □

Consider the variable-radius sPRM algorithm. The following theorem asserts that variable-radius sPRM algorithm is not probabilistically complete in the subcritical regime.

Theorem 20 (Incompleteness of variable-radius sPRM with $r(n) = \gamma n^{-1/d}$) *There exists a constant $\gamma > 0$ such that the variable radius sPRM with connection radius $r(n) = \gamma n^{-1/d}$ is not probabilistically complete.*

The proof of this result requires some intermediate results from random geometric graph theory. Recall that λ_c is the critical density, or continuum percolation threshold (see Section 2.2). Given a Borel set $\Gamma \subseteq \mathbb{R}^d$, let $G_\Gamma^{\text{disc}}(n, r)$ denote the random r-disc graph formed with vertices independent and uniformly sampled from Γ and edges connecting two vertices, v and v', whenever $\|v - v'\| < r_n$.

Lemma 21 (Penrose (2003)) *Let $\lambda \in (0, \lambda_c)$ and $\Gamma \subset \mathbb{R}^d$ be a Borel set. Consider a sequence $\{r_n\}_{n\in\mathbb{N}}$ that satisfies $n\,r_n^d \leq \lambda$, $\forall n \in \mathbb{N}$. Let $N_{\max}(G_\Gamma^{\text{disc}}(n, r_n))$ denote the size of the largest component in $G_\Gamma^{\text{disc}}(n, r_n)$. Then, there exist constants $a, b > 0$ and $m_0 \in \mathbb{N}$ such that for all $m \geq m_0$,*

$$\mathbb{P}\left(\left\{N_{\max}(G_\Gamma^{\text{disc}}(n, r_n)) \geq m\right\}\right) \leq n\left(e^{-a\,m} + e^{-b\,n}\right).$$

Proof of Theorem 20 Let $\epsilon > 0$ such that $\epsilon < \inf_{x \in X_{\text{goal}}} \|x - x_{\text{init}}\|$ and that the $2\,\epsilon$-ball centered at x_{init} lies entirely within the obstacle-free space. Let $G_n^{\text{PRM}} = (V_n^{\text{PRM}}, E_n^{\text{PRM}})$ denote the graph returned by this variable radius sPRM algorithm, when the algorithm is run with n samples. Let $G_n = (V_n, E_n)$ denote the the restriction of G_n^{PRM} to the $2\,\epsilon$-ball centered at x_{init} defined as $V_n = V_n^{\text{PRM}} \cap \mathcal{B}_{x_{\text{init}}, 2\,\epsilon}$ and $E_n = (V_n \times V_n) \cap E_n^{\text{PRM}}$.

Clearly, G_n is equivalent to the random r-disc graph on $\Gamma = \mathcal{B}_{x_{\text{init}}, 2\,\epsilon}$. Let $N_{\max}(G_n)$ denote the number of vertices in the largest connected component of G_n. By Lemma 21, there exists constants $a, b > 0$ and $m_0 \in \mathbb{N}$ such that

$$\mathbb{P}(\{N_{\max}(G_n) \geq m\}) \leq n\left(e^{-a\,m} + e^{-b\,n}\right),$$

for all $m \geq m_0$. Then, for all $m = \lambda^{-1/d}\,(\epsilon/2)\,n^{1/d} > m_0$,

$$\mathbb{P}\left(\left\{N_{\max}(G_n) \geq \lambda^{-1/d}\frac{\epsilon}{2}\,n^{1/d}\right\}\right) \leq n\left(e^{-a\,\lambda^{-1/d}\,(\epsilon/2)\,n^{1/d}} + e^{-b\,n}\right).$$

Let L_n denote the total length of all the edges in the connected component that includes x_{init}. Since $r_n = \lambda^{1/d} n^{-1/d}$,

$$\mathbb{P}\left(\left\{L_n \geq \frac{\epsilon}{2}\right\}\right) \leq n\left(e^{-a\,\lambda^{-1/d}\,(\epsilon/2)\,n^{1/d}} + e^{-b\,n}\right).$$

Since the right hand side is summable, by the Borel-Cantelli lemma the event $\{L_n \geq \epsilon/2\}$ occurs infinitely often with probability zero, i.e., $\mathbb{P}(\limsup_{n\to\infty}\{L_n \geq \epsilon/2\}) = 0$.

Given a graph $G = (V, E)$ define the diameter of this graph as the distance between the farthest pair of vertices in V, i.e., $\max_{v,v' \in V} \|v - v'\|$. Let D_n denote the diameter of the largest component in G_n. Clearly, $D_n \leq L_n$ holds surely. Thus, $\mathbb{P}\left(\limsup_{n\to\infty}\{D_n \geq \epsilon/2\}\right) = 0$.

Let $I \in \mathbb{N}$ be the smallest number that satisfies $r_I \leq \epsilon/2$. Notice that the edges connected to the vertices $V_n^{\mathrm{PRM}} \cap \mathcal{B}_{x_{\mathrm{init}},\epsilon}$ coincide with those connected to $V_n \cap \mathcal{B}_{x_{\mathrm{init}},\epsilon}$, for all $n \geq I$. Let R_n denote distance of the farthest vertex $v \in V_n^{\mathrm{PRM}}$ to x_{init} in the component that contains x_{init} in G_n^{PRM}. Notice also that $R_n \geq \epsilon$ only if $D_n \geq \epsilon/2$, for all $n \geq I$. That is, for all $n \geq I$, $\{R_n \geq \epsilon\} \subseteq \{D_n \geq \epsilon/2\}$, which implies $\mathbb{P}\left(\limsup_{n\to\infty}\{R_n \geq \epsilon\}\right) = 0$.

Let A_n denote the event that the graph returned by this variable radius sPRM algorithm includes a path that reaches the goal region. Clearly, $\{R_n \geq \epsilon\}$ holds, whenever A_n holds. Hence, $\mathbb{P}(A_n) \leq \mathbb{P}(\{R_n \geq \epsilon\})$. Taking the limit superior of both sides yields

$$\liminf_{n\to\infty} \mathbb{P}(A_n) \;\leq\; \limsup_{n\to\infty} \mathbb{P}(A_n) \;\leq\; \limsup_{n\to\infty} \mathbb{P}\left(\{R_n \geq \epsilon\}\right) \;\leq\; \mathbb{P}\Big(\limsup_{n\to\infty}\{R_n \geq \epsilon\}\Big) = 0.$$

Hence, $\lim_{n\to\infty} \mathbb{P}(A_n) = 0$. □

Finally, the probabilistic completeness of the new algorithms proposed in Section 3 is established. Probabilistic completeness of PRM* is implied by its asymptotic optimality, proved in Section 4.2.

Theorem 22 (Completeness of PRM*) *The PRM* algorithm is probabilistically complete.*

Probabilistic completeness of RRG and RRT* is a straightforward consequence of the probabilistic completeness of RRT:

Theorem 23 (Probabilistic completeness of RRG and RRT*) *The RRG and RRT* algorithms are probabilistically complete. Furthermore, for any robustly feasible path planning problem* $(\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$*, there exist constants* $a > 0$ *and* $n_0 \in \mathbb{N}$*, both dependent only on* $\mathcal{X}_{\mathrm{free}}$ *and* $\mathcal{X}_{\mathrm{goal}}$*, such that*

$$\mathbb{P}\left(\left\{V_n^{\mathrm{RRG}} \cap \mathcal{X}_{\mathrm{goal}} \neq \emptyset\right\}\right) > 1 - e^{-a\,n}, \quad \forall n > n_0,$$

and

$$\mathbb{P}\left(\left\{V_n^{\mathrm{RRT}^*} \cap \mathcal{X}_{\mathrm{goal}} \neq \emptyset\right\}\right) > 1 - e^{-a\,n}, \quad \forall n > n_0.$$

Proof By construction, $V_n^{\mathrm{RRG}}(\omega) = V_n^{\mathrm{RRT}^*}(\omega) = V_n^{\mathrm{RRT}}(\omega)$, for all $\omega \in \Omega$ and $n \in \mathbb{N}$. Moreover, the RRG and RRT* algorithms return connected graphs. Hence the result follows directly from the probabilistic completeness of RRT. □

In particular, note that if the RRT algorithm returns a feasible solution by iteration n, so will the RRG and RRT* algorithms, assuming the same sample sequence.

4.2 Asymptotic Optimality

In this section, the optimality problem of path planning is considered. The algorithms presented in Section 3 are analyzed, in terms of their ability to return solutions whose cost converge to the global optimum. First, a definition of asymptotic optimality is provided as almost-sure convergence to optimal paths. Second, it is shown that the RRT algorithm lacks the asymptotic optimality property. Third, the PRM*, RRG, and RRT* algorithms, as well as their k-nearest implementations, are shown to be asymptotically optimal.

Recall from Section 4.1 that an algorithm is probabilistically complete if the algorithm finds with high probability a solution to path planning problems that are robustly feasible, i.e., for which feasible path exists with strong δ-clearance. A similar approach is used to define asymptotic optimality, relying on a notion of weak δ-clearance and on a continuity property for the cost of paths, which will be introduced below.

Let $\sigma_1, \sigma_2 \in \Sigma_{\mathrm{free}}$ be two collision-free paths with the same end points. A path σ_1 is said to be homotopic to σ_2, if there exists a continuous function $\psi : [0,1] \to \Sigma_{\mathrm{free}}$, called the *homotopy*, such that $\psi(0) = \sigma_1$, $\psi(1) = \sigma_2$, and $\psi(\tau)$ is a collision-free path in for all $\tau \in [0,1]$. Intuitively, a path that is homotopic to σ can be continuously transformed to σ through $\mathcal{X}_{\mathrm{free}}$ (see Munkres, 2000). A collision-free path $\sigma : [0,s] \to \mathcal{X}_{\mathrm{free}}$ is said to have *weak δ-clearance*, if there exists a path σ' that has strong δ-clearance and there exist a homotopy ψ, with $\psi(0) = \sigma$, $\psi(1) = \sigma'$, and for all $\alpha \in (0,1]$ there exists $\delta_\alpha > 0$ such that $\psi(\alpha)$ has strong δ_α-clearance. See Figure 2 for an illustration of the weak δ-clearance property. A path that violates the weak δ-clearance property is shown in Figure 3. Weak δ-clearance does not require points along a path to be at least a distance δ away from the obstacles (see Figure 4). In fact, a collision-free path with uncountably many points lying on the boundary of an obstacle can still have weak δ-clearance.

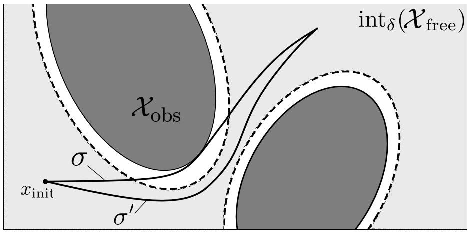
$\mathrm{int}_\delta(\mathcal{X}_{\mathrm{free}})$
$\mathcal{X}_{\mathrm{obs}}$
σ
x_{init}
σ'

Figure 2: An illustration of a path σ with weak δ-clearance. The path σ' that lies inside $\mathrm{int}_\delta(\mathcal{X}_{\mathrm{free}})$ and is in the same homotopy class as σ is also shown in the figure. Note that σ does not have strong δ-clearance.

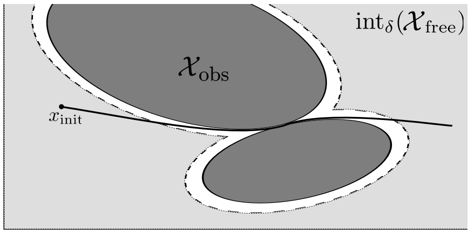
$\mathrm{int}_\delta(\mathcal{X}_{\mathrm{free}})$
$\mathcal{X}_{\mathrm{obs}}$
x_{init}

Figure 3: An illustration of an example path σ that does not have weak δ-clearance. For any positive value of δ, there is no path in $\mathrm{int}_\delta(\mathcal{X}_{\mathrm{free}})$ that is in the same homotopy class as σ.

Next, the set of all paths with bounded length is introduced as a normed space, which allows taking the limit of a sequence of paths. Recall that Σ is the set of all paths, and $TV(\cdot)$ denotes the

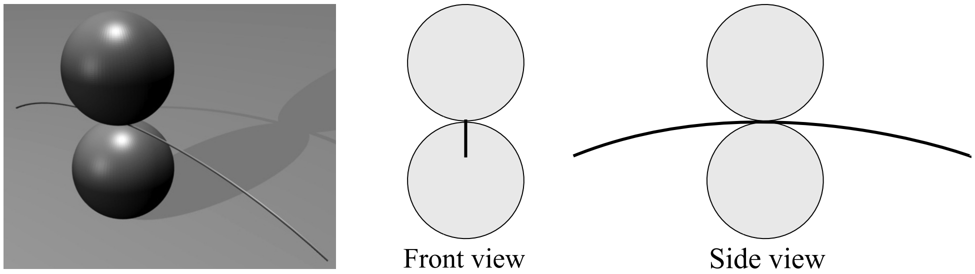

Figure 4: An illustration of a path that has weak δ-clearance. The path passes through a point where two spheres representing the obstacle region are in contact. Clearly, the path does not have strong δ-clearance.

total variation, i.e., the length, of a path (see Section 2.1). Given $\sigma_1, \sigma_2 \in \Sigma$ with $\sigma_1 : [0,1] \to \mathcal{X}$ and $\sigma_2 : [0,1] \to \mathcal{X}$, the addition operation is defined as $(\sigma_1 + \sigma_2)(\tau) = \sigma_1(\tau) + \sigma_2(\tau)$ for all $\tau \in [0,1]$. The set of paths Σ is closed under addition. Given a path $\sigma : [0,1] \to \mathcal{X}$ and a scalar $\alpha \in \mathbb{R}$, the multiplication by a scalar operation is defined as $(\alpha\sigma)(\tau) := \alpha\, \sigma(\tau)$ for all $\tau \in [0,1]$. With these addition and multiplication by a scalar operations, the function space Σ is, in fact, a vector space. On the vector space Σ, define the norm $\|\sigma\|_{\mathrm{BV}} := \int_0^1 |\sigma(\tau)|\, d\tau + \mathrm{TV}(\sigma)$, and denote the function space Σ endowed with the norm $\|\cdot\|_{\mathrm{BV}}$ by $\mathrm{BV}(\mathcal{X})$. The norm $\|\cdot\|_{\mathrm{BV}}$ induces the following distance function:

$$\mathrm{dist}(\sigma_1, \sigma_2) = \|\sigma_1 - \sigma_2\|_{\mathrm{BV}} = \int_0^1 \big\|(\sigma_1 - \sigma_2)(\tau)\big\| d\tau + \mathrm{TV}(\sigma_1 - \sigma_2)$$

where $\|\cdot\|$ is the usual Euclidean norm. A sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths is said to converge to a path $\bar{\sigma}$, denoted as $\lim_{n\to\infty} \sigma_n = \bar{\sigma}$, if the norm of the difference between σ_n and $\bar{\sigma}$ converges to zero, i.e., $\lim_{n\to\infty} \|\sigma_n - \bar{\sigma}\|_{\mathrm{BV}} = 0$.

A feasible path $\sigma^* \in \mathcal{X}_{\mathrm{free}}$ that solves the optimality problem (Problem 3) is said to be a *robustly optimal solution* if it has weak δ-clearance and, for any sequence of collision-free paths $\{\sigma_n\}_{n\in\mathbb{N}}$, $\sigma_n \in \mathcal{X}_{\mathrm{free}}$, $\forall n \in \mathbb{N}$, such that $\lim_{n\to\infty} \sigma_n = \sigma^*$, $\lim_{n\to\infty} c(\sigma_n) = c(\sigma^*)$. Clearly, a path planning problem that has a robustly optimal solution is necessarily robustly feasible. Let $c^* = c(\sigma^*)$ be the cost of an optimal path, and let Y_n^{ALG} be the extended random variable corresponding to the cost of the minimum-cost solution included in the graph returned by ALG at the end of iteration n.

Definition 24 (Asymptotic Optimality) *An algorithm ALG is asymptotically optimal if, for any path planning problem $(\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$ and cost function $c : \Sigma \to \mathbb{R}_{\geq 0}$ that admit a robustly optimal solution with finite cost c^*,*

$$\mathbb{P}\left(\left\{\limsup_{n\to\infty} Y_n^{\mathrm{ALG}} = c^*\right\}\right) = 1.$$

Note that, since $Y_n^{\mathrm{ALG}} \geq c^*$, $\forall n \in \mathbb{N}$, asymptotic optimality of ALG implies that the limit $\lim_{n\to\infty} Y_n^{\mathrm{ALG}}$ exists, and is equal to c^*. Clearly, probabilistic completeness is necessary for asymptotic optimality. Moreover, the probability that a sampling-based algorithm converges to an optimal solution almost surely has probability either zero or one. That is, a sampling-based algorithm either converges to the optimal solution in almost all runs, or the convergence does not occur in almost all runs.

Lemma 25 *Given that* $\limsup_{n\to\infty} Y_n^{\mathrm{ALG}} < \infty$*, i.e.,* ALG *finds a feasible solution eventually, the probability that* $\limsup_{n\to\infty} Y_n^{\mathrm{ALG}} = c^*$ *is either zero or one.*

Proof Conditioning on the event $\{\limsup_{n\to\infty} Y_n^{\mathrm{ALG}} < \infty\}$ ensures that Y_n^{ALG} is finite, thus a random variable, for all large n. Given a sequence $\{Y_n\}_{n\in\mathbb{N}}$ of random variables, let $\mathcal{F}'_m$ denote the σ-field generated by the sequence $\{Y_n\}_{n=m}^{\infty}$ of random variables. The tail σ-field $\mathcal{T}$ is defined as $\mathcal{T} = \bigcap_{n\in\mathbb{N}} \mathcal{F}'_n$. An event A is said to be a *tail event* if $A \in \mathcal{T}$. Any tail event occurs with probability either zero or one by the Kolmogorov zero-one law (Resnick, 1999). Consider the sequence $\{Y_n^{\mathrm{ALG}}\}_{n\in\mathbb{N}}$ of random variables. Let $\mathcal{F}'_m$ denote the σ-fields generated by $\{Y_n^{\mathrm{ALG}}\}_{n=m}^{\infty}$. Then, $\left\{\limsup_{n\to\infty} Y_n^{\mathrm{ALG}} = c^*\right\} = \left\{\limsup_{n\to\infty,\, n\geq m} Y_n^{\mathrm{ALG}} = c^*\right\} \in \mathcal{F}'_m$ for all $n \in \mathbb{N}$. Hence, $\left\{Y_n^{\mathrm{ALG}} = c^*\right\} \in \bigcap_{n\in\mathbb{N}} \mathcal{F}'_m$ is a tail event. The result follows by the Kolmogorov zero-one law. $\square$

Among the first steps in assessing the asymptotic optimality properties of an algorithm ALG is determining whether the limit $\lim_{n\to\infty} Y_n^{\mathrm{ALG}}$ exists. It turns out that if the graphs returned by ALG satisfy a monotonicity property, then the limit exists, and is in general a random variable, indicated with Y_∞^{ALG}.

Lemma 26 *If* $G_i^{\mathrm{ALG}}(\omega) \subseteq G_{i+1}^{\mathrm{ALG}}(\omega)$, $\forall \omega \in \Omega$ *and* $\forall i \in \mathbb{N}$, *then* $\lim_{n\to\infty} Y_n^{\mathrm{ALG}}(\omega) = Y_\infty^{\mathrm{ALG}}(\omega)$.

Proof Since $G_i^{\mathrm{ALG}}(\omega) \subseteq G_{i+1}^{\mathrm{ALG}}(\omega)$, then $Y_{i+1}^{\mathrm{ALG}}(\omega) \leq Y_i^{\mathrm{ALG}}(\omega)$, for all $\omega \in \Omega$. Since $Y_i^{\mathrm{ALG}} \geq c^*$, then the sequence converges to some limiting value, dependent on ω, i.e., $Y_\infty^{\mathrm{ALG}}(\omega)$. $\square$

Of the algorithms presented in Section 3, it is easy to check that PRM, sPRM, RRT, RRG, and RRT* satisfy the monotonicity property in Lemma 26. On the other hand, k-nearest sPRM and PRM* do not: in these cases, the random variable Y_{i+1}^{ALG} is not necessarily dominated by Y_i^{ALG}. This is evident in numerical experiments, e.g., see Figures 10 and 11 in Section 5.

In order to avoid trivial cases of asymptotic optimality, it is necessary to rule out problems in which optimal solutions can be computed after a finite number of samples. Let Σ^* denote the set of all optimal paths, i.e., the set of all paths that solve the optimal planning problem (Problem 3), and $\mathcal{X}_{\mathrm{opt}}$ denote the set of states that an optimal path in Σ^* passes through, i.e.,

$$\mathcal{X}_{\mathrm{opt}} = \{x \in X_{\mathrm{free}} \,|\, \exists \sigma^* \in \Sigma^*, \tau \in [0,1] \text{ such that } x = \sigma^*(\tau)\}.$$

Assumption 27 (Zero-measure Optimal Paths) *The set of all points traversed by an optimal trajectory has measure zero, i.e.,* $\mu(\mathcal{X}_{\mathrm{opt}}) = 0$.

Most cost functions and problem instances of interest satisfy this assumption, including, e.g., the Euclidean length of the path when the goal region is convex. This assumption does not imply that there is a single optimal path; indeed, there are problem instances with uncountably many optimal paths, for which Assumption 27 holds. (A simple example is the motion planning problem in three dimensional Euclidean space where a ball shaped obstacle is placed between the initial state and the goal region.) Assumption 27 implies that no sampling-based planning algorithm can find a solution to the optimality problem in a finite number of iterations.

Lemma 28 *If Assumption 27 holds, the probability that a sampling-based algorithm* ALG *returns a graph containing an optimal path at a finite iteration* $n \in \mathbb{N}$ *is zero, i.e.,*

$$\mathbb{P}\left(\cup_{n\in\mathbb{N}} \{Y_n^{\mathrm{ALG}} = c^*\}\right) = 0.$$

Proof Let B_n denote the event that ALG constructs a graph containing a path with cost exactly equal to c^* at the end of iteration i, i.e., $B_n = \{Y_n^{\mathrm{ALG}} = c^*\}$. Let B denote the event that ALG returns a graph containing a path that costs exactly c^* at some finite iteration i. Then, B can be written as $B = \cup_{n\in\mathbb{N}} B_n$. Since $B_n \subseteq B_{n+1}$, by monotonocity of measures, $\lim_{i\to\infty} \mathbb{P}(B_n) = \mathbb{P}(B)$. By Assumption 27 and the definition of the sampling procedure, $\mathbb{P}(B_n) = 0$ for all $n \in \mathbb{N}$, since the probability that the set $\bigcup_{i=1}^{n} \{\texttt{SampleFree}(i)\}$ of points contains a point from a zero-measure set is zero. Hence, $\mathbb{P}(B) = 0$. □

In the remainder of the paper, it will be tacitly assumed that Assumption 27, and hence Lemma 28, hold.

4.2.1 Existing algorithms

The algorithms in Section 3.2 were originally introduced to efficiently solve the feasibility problem, relaxing the completeness requirement to probabilistic completeness. Nevertheless, it is of interest to establish whether these algorithms are asymptotically optimal in addition to being probabilistically complete. (The first two results in this section rely on results that will be proven in Section 4.2.2, i.e., the fact that the RRT algorithm is not asymptotically optimal, and the PRM* algorithm is asymptotically optimal)

First, consider the PRM algorithm and its variants. The PRM algorithm, in its original form, is not asymptotically optimal.

Theorem 29 (Non-optimality of PRM) *The PRM algorithm is not asymptotically optimal.*

Proof The proof is based on a counterexample, establishing a form of equivalence between PRM and RRT, which in turn will be proven not to be asymptotically optimal in Theorem 33. Consider a convex obstacle-free environment, e.g., $\mathcal{X}_{\mathrm{free}} = \mathcal{X}$, and choose the connection radius for PRM and the steering parameter for RRT such that $r, \eta > \mathrm{diam}(\mathcal{X})$. At each iteration, exactly one vertex and one edge is added to the graph, since (i) all connection attempts using the local planner (e.g., straight line connections as considered in this paper) are collision-free, and (ii) at the end of each iteration, the graph is connected (i.e., it contains only one connected component). In particular, the graph returned by the PRM algorithm in this case is a tree, and the arborescence obtained by choosing as the root the first sample point, i.e., $\texttt{SampleFree}_0$, is an online nearest-neighbor graph (see Section 2.2) coinciding with the graph returned by RRT with the random initial condition $x_{\mathrm{init}} = \texttt{SampleFree}_0$.

Recall that the PRM algorithm is applicable for multiple-query planning problems: in other words, the graph returned by the PRM algorithm is used to solve path planning problems from arbitrary $x_{\mathrm{init}} \in \mathcal{X}_{\mathrm{free}}$ and $\mathcal{X}_{\mathrm{goal}} \subset \mathcal{X}_{\mathrm{free}}$. (Note that all such problems admit robust optimal solutions.) In particular, for $x_{\mathrm{init}} = \texttt{SampleFree}_0$, and any X_{goal}, then $Y_n^{\mathrm{PRM}}(\omega) = Y_n^{\mathrm{RRT}}(\omega)$, for all $\omega \in \Omega$, $n \in \mathbb{N}$. In particular, since both PRM and RRT satisfy the monotonicity condition in Lemma 26, Theorem 33 implies that

$$\mathbb{P}\left(\left\{\limsup_{n\to\infty} Y_n^{\mathrm{PRM}} = c^*\right\}\right) = \mathbb{P}\left(\left\{\lim_{n\to\infty} Y_n^{\mathrm{PRM}} = c^*\right\}\right) = \mathbb{P}\left(\left\{\lim_{n\to\infty} Y_n^{\mathrm{RRT}} = c^*\right\}\right) = 0.$$

□

The lack of asymptotic optimality of PRM is due to its incremental construction, coupled with the constraint eliminating edges making unnecessary connections within a connected component. Such a constraint is not present in the batch construction of the sPRM algorithm, which is indeed asymptotically optimal (at the expense of computational complexity, see Section 4.3).

Theorem 30 (Asymptotic Optimality of sPRM) *The sPRM algorithm is asymptotically optimal.*

Proof By construction, $V_n^{\mathrm{sPRM}}(\omega) = V_n^{\mathrm{PRM}^*}(\omega)$, and $E_n^{\mathrm{sPRM}}(\omega) \supseteq E_n^{\mathrm{PRM}^*}(\omega)$ for all $\omega \in \Omega$. Hence, the graph returned by sPRM includes all the paths that are present in the graph returned by PRM*. Then, asymptotic optimality of sPRM follows from that of PRM*, which will be proven in Theorem 34. ☐

On the other hand, as in the case of probabilistic completeness, the heuristics that are often used in the practical implementation of (s)PRM are not asymptotically optimal.

Theorem 31 (Non-optimality of k-nearest sPRM) *The k-nearest sPRM algorithm is not asymptotically optimal, for any constant $k \in \mathbb{N}$.*

This theorem will be proven under the assumption that the underlying point process is Poisson. More precisely, the algorithm is analyzed when it is run with Poisson(n) samples. That is, the realization of the random variable Poisson(n) determines the number of points sampled independently and uniformly in $\mathcal{X}_{\mathrm{free}}$. Hence, the expected number of samples is equal to n, although its realization may slightly differ. However, since the Poisson random variable has exponentially-decaying tails, its large deviations from its mean is unlikely (see, e.g., Grimmett and Stirzaker (2001) for a more precise statement). With a slight abuse of notation, the cost of the best path in the graph returned by the k-nearest sPRM algorithm when the algorithm is run with Poisson(n) number of samples is denoted by $Y_n^{k\mathrm{PRM}}$, and it is shown that $\mathbb{P}(\{\limsup_{n\to\infty} Y_n^{k\mathrm{PRM}} = c^*\}) = 0$.

Proof of Theorem 31 Let σ^* denote an optimal path and s^* denote its length, i.e., $s^* = TV(\sigma^*)$. For each n, consider a tiling of σ^* with disjoint open hypercubes, each with edge length $2\, n^{-1/d}$, such that the center of each cube is a point on σ^*. See Figure 5. Let M_n denote the maximum number of tiles that can be generated in this manner and note $M_n \geq \frac{s^*}{2}\, n^{1/d}$. Partition each tile into several open cubes as follows: place an inner cube with edge length $n^{-1/d}$ at the center of the tile and place several outer cubes each with edge length $\frac{1}{2}\, n^{-1/d}$ around the cube at the center as shown in Figure 5. Let F_d denote the number of outer cubes. The volumes of the inner cube and each of the outer cubes are n^{-1} and $2^{-d}\, n^{-1}$, respectively.

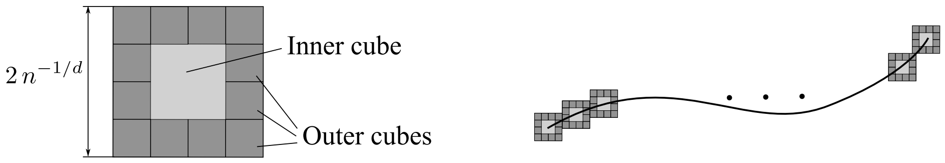

Figure 5: An illustration of the tiles mention in the proof of Theorem 31. A single tile is shown in the left; a tiling of the optimal trajectory σ^* is shown on the right.

For $n \in \mathbb{N}$ and $m \in \{1, 2, \ldots, M_n\}$, consider the tile m when the algorithm is run with Poisson(n) samples. Let $I_{n,m}$ denote the indicator random variable for the event that the center cube of this tile contains no samples, whereas every outer cube contains at least $k+1$ samples, in tile m.

The probability that the inner cube contains no samples is $e^{-1/\mu(\mathcal{X}_{\mathrm{free}})}$. The probability that an outer cube contains at least $k+1$ samples is $1 - \mathbb{P}\left(\{\mathrm{Poisson}(2^{-d}/\mu(\mathcal{X}_{\mathrm{free}})) \geq k+1\}\right) = 1 - \mathbb{P}(\{\mathrm{Poisson}(2^{-d}/\mu(\mathcal{X}_{\mathrm{free}})) \leq k\}) = 1 - \frac{\Gamma(k+1, 2^{-d}/\mu(\mathcal{X}_{\mathrm{free}}))}{k!}$, where $\Gamma(\cdot,\cdot)$ is the incomplete gamma function (Abramowitz and Stegun, 1964). Then, noting that the cubes in a given tile are disjoint and using the independence property of the Poisson process (see Lemma 11),

$$\mathbb{E}\left[I_{n,m}\right] \;=\; e^{-1/\mu(\mathcal{X}_{\text{free}})} \left(1 - \frac{\Gamma(k+1, 2^{-d}/\mu(\mathcal{X}_{\text{free}}))}{k!}\right)^{F_d} \;>\; 0,$$

which is a constant that is independent of n; denote this constant by α.

Let $G_n = (V_n, E_n)$ denote the graph returned by the k-nearest PRM algorithm by the end of Poisson(n) iterations. Observe that if $I_{n,m} = 1$, then there is no edge of G_n crossing the cube of side length $\frac{1}{2}\, n^{-1/d}$ that is centered at the center of the inner cube in tile m (shown as the white cube in Figure 6). To prove this claim, note the following two facts. First, no point that is outside of the cubes can have an edge that crosses the inner cube. Second, no point in one of the outer cubes has an edge that has length greater than $\frac{\sqrt{d}}{2}\, i^{-1/d}$. Thus, no edge can cross the white cube illustrated in Figure 6.

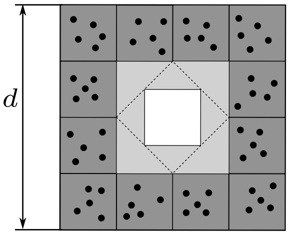

Figure 6: The event that the inner cube contains no points and each outer cube contains at least k points of the point process is illustrated. The cube of side length $\frac{1}{2}\, n^{-1/d}$ is shown in white.

Let σ_n denote the path in G_n that is closest to σ^* in terms of the bounded variation norm. Let $U_n := \|\sigma_n - \sigma^*\|_{\text{BV}}$. Notice that $U_n \geq \frac{1}{2}\, n^{-1/d} \sum_{m=1}^{M_n} I_{n,m} = \frac{1}{2}\, n^{-1/d}\, M_n\, I_{n,1} = \frac{s^*}{4} I_{n,1}$. Then,

$$\mathbb{E}\left[\limsup_{n\to\infty} U_n\right] \;\geq\; \limsup_{n\to\infty} \mathbb{E}\left[U_n\right] \;\geq\; \limsup_{n\to\infty} \frac{s^*}{4}\, \mathbb{E}\left[I_{n,m}\right] \;\geq\; \frac{\alpha\, s^*}{4} \;>\; 0,$$

where the first inequality follows from Fatou's lemma (Resnick, 1999). This implies $\mathbb{P}(\{\limsup_{n\to\infty} U_n > 0\}) > 0$. Since $U_i > 0$ implies $Y_n > c^*$ surely,

$$\mathbb{P}\left(\{\limsup_{n\to\infty} Y_n > c^*\}\right) \geq \mathbb{P}\left(\{\limsup_{n\to\infty} U_n > 0\}\right) > 0.$$

That is, $\mathbb{P}\left(\{\limsup_{n\to\infty} Y_n = c^*\}\right) < 1$. In fact, by Lemma 25, $\mathbb{P}\left(\{\limsup_{n\to\infty} Y_n = c^*\}\right) = 0$. $\square$

Second, asymptotic optimality of a large class of variable radius sPRM algorithms is considered. Consider a variable radius sPRM in which connection radius satisfies $r(n) \leq \gamma\, n^{-1/d}$ for some $\gamma > 0$ and for all $n \in \mathbb{N}$. The next theorem shows that this algorithm lacks the asymptotic optimality property.

Theorem 32 (Non-optimality of variable radius sPRM with $r(n) = \gamma\, n^{-1/d}$) *Consider a variable radius sPRM algorithm with connection radius $r(n) = \gamma\, n^{-1/d}$. This sPRM algorithm is not asymptotic optimal for any $\gamma \in \mathbb{R}_{\geq 0}$.*

Proof Let σ^* denote a path that is a robust solution to the optimality problem. Let n denote the number of samples that the algorithm is run with. For all n, construct a set $B_n = \{B_{n,1}, B_{n,2}, \ldots, B_{n,M_n}\}$ of openly disjoint balls as follows. Each ball in B_n has radius $r_n = \gamma\, n^{-1/d}$, and lies entirely inside $\mathcal{X}_{\rm free}$. Furthermore, the balls in B_n "tile" σ^* such that the center of each ball lies on σ^* (see Figure 7). Let M_n denote the maximum number of balls, $\bar{s}$ denote the length of the portion of σ^* that lies within the δ-interior of $\mathcal{X}_{\rm free}$, and $n_0 \in \mathbb{N}$ denote the number for which $r_n \leq \delta$ for all $n \geq n_0$.

Then, for all $n \geq n_0$,

$$M_n \geq \frac{\bar{s}}{2\,\gamma\,\left(\frac{1}{n}\right)^{1/d}} = \frac{\bar{s}}{2\,\gamma}\, n^{1/d}.$$

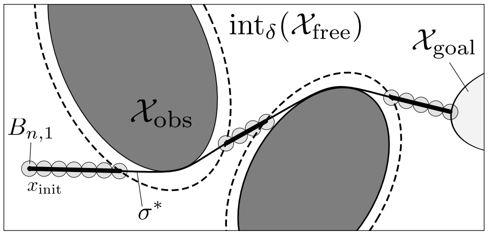

Figure 7: An illustration of the covering of the optimal path, σ^*, with openly disjoint balls. The balls cover only a portion of σ^* that lies within the δ-interior of $\mathcal{X}_{\rm free}$.

Indicate the graph returned by this sPRM algorithm as $G_n = (V_n, E_n)$. Denote the event that the ball $B_{n,m}$ contains no vertex in V_n by $A_{n,m}$. Denote the indicator random variable for the event $A_{n,m}$ by $I_{n,m}$, i.e., $I_{n,m} = 1$ when $A_{n,m}$ holds and $I_{n,m} = 0$ otherwise. Then, for all $n \geq n_0$,

$$\mathbb{E}[I_{n,m}] = \mathbb{P}(A_{n,m}) = \left(1 - \frac{\mu(B_{n,m})}{\mu(\mathcal{X}_{\rm free})}\right)^n = \left(1 - \frac{\zeta_d\,\gamma^d}{\mu(\mathcal{X}_{\rm free})}\,\frac{1}{n}\right)^n$$

Let N_n be the random variable that denotes the total number of balls in B_n that contain no vertex in V_n, i.e., $N_n = \sum_{m=1}^{M_n} I_{n,m}$. Then, for all $n \geq n_0$,

$$\mathbb{E}[N_n] \;=\; \mathbb{E}\left[\sum\nolimits_{m=1}^{M_n} I_{n,m}\right] \;=\; \sum_{m=1}^{M_n} \mathbb{E}[I_{n,m}] \;=\; M_n\; \mathbb{E}[I_{n,1}] \;\geq\; \frac{\bar{s}}{2\,\gamma}\, n^{1/d} \left(1 - \frac{\zeta_d\,\gamma^d}{\mu(\mathcal{X}_{\rm free})}\,\frac{1}{n}\right)^n.$$

Consider a ball $B_{n,m}$ that contains no vertices of this sPRM algorithm. Then, no edges of the graph returned by this algorithm cross the ball of radius $\frac{\sqrt{3}}{2} r_n$ centered at the center of $B_{n,m}$. See Figure 8.

Let P_n denote the (finite) set of all acyclic paths that reach the goal region in the graph returned by this sPRM algorithm when the algorithm is run with n samples. Let U_n denote the total variation of the path that is closest to σ^* among all paths in P_n, i.e., $U_n := \min_{\sigma_n \in P_n} \|\sigma_n - \sigma^*\|_{\rm BV}$. Then,

$$\mathbb{E}[U_n] \;\geq\; \mathbb{E}\left[\gamma \left(\frac{1}{n}\right)^{1/d} N_n\right] \;\geq\; \frac{\bar{s}}{2} \left(1 - \frac{\zeta_d\,\gamma^d}{\mu(\mathcal{X}_{\rm free})}\frac{1}{n}\right)^n.$$

Taking the limit superior of both sides, the following inequality can be established:

$$\mathbb{E}\left[\limsup_{n\to\infty} U_n\right] \;\geq\; \limsup_{n\to\infty} \mathbb{E}\left[U_n\right] \;\geq\; \limsup_{n\to\infty} \frac{\bar{s}}{2} \left(1 - \frac{\zeta_d\,\gamma^d}{\mu(\mathcal{X}_{\rm free})}\frac{1}{n}\right)^n \;=\; \frac{\bar{s}}{2}\, e^{-\frac{\zeta_d\,\gamma^d}{\mu(\mathcal{X}_{\rm free})}} > 0,$$

where the first inequality follows from Fatou's lemma (Resnick, 1999). Hence, $\mathbb{P}(\{\limsup_{n\to\infty} U_n > 0\}) > 0$, which implies that $\mathbb{P}\left(\left\{\limsup_{n\to\infty} Y_n^{\rm ALG} > c^*\right\}\right) > 0$. That is, $\mathbb{P}\left(\left\{\limsup_{n\to\infty} Y_n^{\rm ALG} = c^*\right\}\right) < 1$. In fact, $\mathbb{P}\left(\left\{\limsup_{n\to\infty} Y_n^{\rm ALG} = c^*\right\}\right) = 0$ by the Kolmogorov zero-one law (see Lemma 25). $\square$

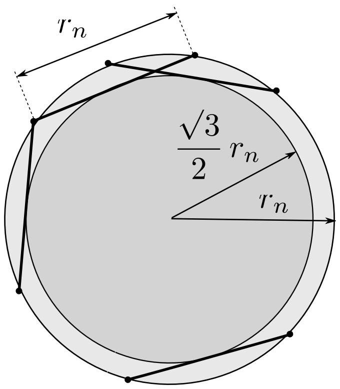

Figure 8: If the outer ball does not contain vertices of the PRM graph, then no edge of the graph corresponds to a path crossing the inner ball.

Rapidly-exploring Random Trees In this section, it is shown that the minimum-cost path in the RRT algorithm converges to a certain random variable, however, under mild technical assumptions, this random variable is *not* equal to the optimal cost, with probability one.

Theorem 33 (Non-optimality of RRT) *The RRT algorithm is not asymptotically optimal.*

The proof of this theorem can be found in Appendix B. Note that, since at each iteration the RRT algorithm either adds a vertex and an edge, or leaves the graph unchanged, $G_i^{\mathrm{RRT}}(\omega) \subseteq G_{i+1}^{\mathrm{RRT}}(\omega)$, for all $i \in \mathbb{N}$ and all $\omega \in \Omega$, and hence the limit $\lim_{n\to\infty} Y_n^{\mathrm{RRT}}$ exists and is equal to the random variable Y_∞^{RRT}. In conjunction with Lemma 25, Theorem 33 implies that this limit is strictly greater than c^* almost surely, i.e., $\mathbb{P}\left(\{\lim_{n\to\infty} Y_n^{\mathrm{RRT}} > c^*\}\right) = 1$. In other words, the cost of the best solution returned by RRT converges to a suboptimal value, with probability one. In fact, it is possible to construct problem instances such that the probability that the first solution returned by the RRT algorithm has arbitrarily high cost is bounded away from zero (Nechushtan et al., 2010).

Since the cost of the best path returned by the RRT algorithm converges to a random variable, Theorem 33 provides new insight explaining the effectiveness of approaches as in Ferguson and Stentz (2006). In fact, running multiple instances of the RRT algorithm amounts to drawing multiple samples of Y_∞^{RRT}.

4.2.2 Proposed algorithms

In this section, the proposed algorithms are analyzed for asymptotic optimality, i.e., almost sure convergence to optimal solutions. It is shown that the PRM*, RRG, and RRT* algorithms, as well as their k-nearest implementations, are all asymptotically optimal. The proofs of the following theorems are quite lengthy, and will be provided in the appendix.

Recall that d denotes the dimensionality of the configuration space, $\mu(\mathcal{X}_{\mathrm{free}})$ denotes the Lebesgue measure of the obstacle-free space, and ζ_d denotes the volume of the unit ball in the d-dimensional Euclidean space. Proofs of the following theorems can be found in Appendices C–G.

Theorem 34 (Asymptotic optimality of PRM*) *If* $\gamma_{\mathrm{PRM}} > 2\,(1 + 1/d)^{1/d} \left(\frac{\mu(X_{\mathrm{free}})}{\zeta_d}\right)^{1/d}$, *then the PRM* algorithm is asymptotically optimal.*

Theorem 35 (Asymptotic optimality of k-nearest PRM*) *If* $k_{\mathrm{PRM}} > e\,(1 + 1/d)$, *then the k-nearest implementation of the PRM* algorithm is asymptotically optimal.*

Theorem 36 (Asymptotic optimality of RRG) *If* $\gamma_{\mathrm{PRM}} > 2\,(1+1/d)^{1/d} \left(\frac{\mu(X_{\mathrm{free}})}{\zeta_d}\right)^{1/d}$*, then the RRG algorithm is asymptotically optimal.*

Theorem 37 (Asymptotic optimality of k-nearest RRG) *If* $k_{\mathrm{RRG}} > e\,(1+1/d)$*, then the k-nearest implementation of the RRG algorithm is asymptotically optimal.*

Theorem 38 (Asymptotic optimality of RRT*) *If* $\gamma_{\mathrm{RRT}^*} > (2\,(1{+}1/d))^{1/d} \left(\frac{\mu(X_{\mathrm{free}})}{\zeta_d}\right)^{1/d}$*, then the RRT* algorithm is asymptotically optimal.*

Theorem 39 (Asymptotic optimality of k-nearest RRT*) *If* $k_{\mathrm{RRT}^*} > 2^{d+1}\,e\,(1+1/d)$*, then the k-nearest implementation of the RRT* algorithm is asymptotically optimal.*

The proof of the latter theorem follows from those of Theorems 37 and 38.

4.3 Computational Complexity

The objective of this section is to compare the computational complexity of the algorithms provided in Section 3. First, each algorithm is analyzed in terms of the number of calls to the `CollisionFree` procedure. Second, the computational complexity of certain primitive procedures such as `Nearest` and `Near` (see Section 3.1) are analyzed. Using these results, a thorough analysis of the computational complexity of the all the algorithms is given in terms of the number of simple operations, such as comparisons, additions, multiplications. An analysis of the computational complexity of the query phase, i.e., the complexity of extracting the optimal solution from the graph returned by these algorithms, is also provided.

The following notation for asymptotic computational complexity will be used throughout this section. Let $W_n^{\mathrm{ALG}}(P)$ be a function of the graph returned by algorithm ALG when ALG is run with inputs $P = (\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$ and n. Clearly, $W_n^{\mathrm{ALG}}(P)$ is a random variable. Let $f : \mathbb{N} \to \mathbb{N}$ be an increasing function with $\lim_{n\to\infty} f(n) = \infty$. The random variable W_n^{ALG} is said belong to $\Omega(f(n))$, denoted as $W_n^{\mathrm{ALG}} \in \Omega(f(n))$, if there exists a problem instance $P = (\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$ such that $\liminf_{n\to\infty} \mathbb{E}[W_n^{\mathrm{ALG}}(P)/f(n)] > 0$. Similarly, W_n^{ALG} is said to belong to $O(f(n))$ if $\limsup_{n\to\infty} \mathbb{E}[W_n^{\mathrm{ALG}}(P)/f(n)] < \infty$ for all problem instances $P = (\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$.

Number of calls to the `CollisionFree` procedure Let M_n^{ALG} denote the total number of calls to the `CollisionFree` procedure by algorithm ALG in iteration n.

First, lower-bounds are established for the PRM and sPRM algorithms.

Lemma 40 (PRM) $M_n^{\mathrm{PRM}} \in \Omega(n)$.

Proof Consider the problem instance $(\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$, where $\mathcal{X}_{\mathrm{free}}$ is composed of two openly-disjoint sets $\mathcal{X}_1$ and $\mathcal{X}_2$ (see Figure 9). The set $\mathcal{X}_2$ is designed to be a hyperrectangle shaped set with one side equal to $r/2$, where r is the connection radius.

Any r-ball centered at a point in $\mathcal{X}_2$ will certainly contain a nonzero measure part of $\mathcal{X}_2$. Define $\bar{\mu}$ as the volume of the smallest region in $\mathcal{X}_2$ that can be intersected by an r-ball centered at $\mathcal{X}_2$, i.e., $\bar{\mu} := \inf_{x\in\mathcal{X}_2} \mu(\mathcal{B}_{x,r} \cap \mathcal{X}_1)$. Clearly, $\bar{\mu} > 0$.

Thus, for any sample X_n that falls into $\mathcal{X}_2$, the PRM algorithm will attempt to connect X_n to a certain number of vertices that lies in a subset $\mathcal{X}_1'$ of $\mathcal{X}_1$ such that $\mu(\mathcal{X}_1') \geq \bar{\mu}$. The expected number of vertices in $\mathcal{X}_1'$ is at least $\bar{\mu}\,n$. Moreover, none of these vertices can be in the same connected component with X_n. Thus, $\mathbb{E}[M_n^{\mathrm{PRM}}/n] > \bar{\mu}$. The result is obtained by taking the limit inferior of both sides. $\square$

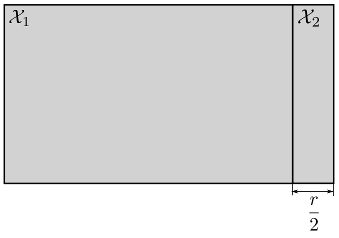

Figure 9: An illustration of $\mathcal{X}_{\mathrm{free}} = \mathcal{X}_1 \cup \mathcal{X}_2$.

Lemma 41 (sPRM) $M_n^{\mathrm{sPRM}} \in \Omega(n)$.

Proof The proof of a stronger result is provided. It is shown that for all problem instances $P = (\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$, $\liminf_{n\to\infty} \mathbb{E}[M_n^{\mathrm{sPRM}}/n] > 0$, which implies the lemma. Recall from Algorithm 2 that r denotes the connection radius. Let $\bar{\mu}$ denote the volume of the smallest region that can be formed by intersecting $\mathcal{X}_{\mathrm{free}}$ with an r-ball centered at a point inside $\mathcal{X}_{\mathrm{free}}$, i.e., $\bar{\mu} := \inf_{x\in\mathcal{X}_{\mathrm{free}}} \mu(\mathcal{B}_{x,r} \cap \mathcal{X}_{\mathrm{free}})$. Recall that $\mathcal{X}_{\mathrm{free}}$ is the closure of an open set. Hence, $\bar{\mu} > 0$.

Clearly, M_n, the number of calls to the `CollisionFree` procedure in iteration n, is equal to the number of nodes inside the ball of radius r centered at the last sample point X_n. Moreover, the volume of the $\mathcal{X}_{\mathrm{free}}$ that lies inside this ball is at least $\bar{\mu}$. Then, the expected value of M_n is lower bounded by the expected value of a binomial random variable with parameters $\bar{\mu}/\mu(\mathcal{X}_{\mathrm{free}})$ and n, since the underlying point process is binomial. Thus, $\mathbb{E}[M_n^{\mathrm{sPRM}}] \geq \frac{\bar{\mu}}{\mu(\mathcal{X}_{\mathrm{free}})}\, n$. Then, $\mathbb{E}[M_n/n] \geq \bar{\mu}/\mathcal{X}_{\mathrm{free}}$ for all $n \in \mathbb{N}$. Taking the limit inferior of both sides gives the result. ☐

Clearly, for k-nearest PRM, $M_n^{k\text{-sPRM}} = k$ for all $n \in \mathbb{N}$ with $n > k$. Similarly, for the RRT, $M_n^{\mathrm{RRT}} = 1$ for all $n \in \mathbb{N}$.

The next lemma upper-bounds the number of calls to the `CollisionFree` procedure in the proposed algorithms.

Lemma 42 (PRM*, RRG, and RRT*) $M_n^{\mathrm{PRM}^*}$, M_n^{RRG}, $M_n^{\mathrm{RRT}^*} \in O(\log n)$.

Proof First, consider PRM*. Recall that r_n denotes the connection radius of the PRM* algorithm. Recall that the r_n interior of $\mathcal{X}_{\mathrm{free}}$, denoted by $\mathrm{int}_{r_n}(\mathcal{X}_{\mathrm{free}})$, is defined as the set of all points x, for which the r_n-ball centered at x lies entirely inside $\mathcal{X}_{\mathrm{free}}$. Let A denote the event that the sample X_n drawn at the last iteration falls into the r_n interior of $\mathcal{X}_{\mathrm{free}}$. Then,

$$\mathbb{E}\big[M_n^{\mathrm{PRM}^*}\big] = \mathbb{E}\big[M_n^{\mathrm{PRM}^*} \,\big|\, A\big]\, \mathbb{P}(A) + \mathbb{E}\big[M_n^{\mathrm{PRM}^*} \,\big|\, A^c\big]\, \mathbb{P}(A^c).$$

Let $n_0 \in \mathbb{N}$ be the smallest number such that $\mu(\mathrm{int}_{r_n}(\mathcal{X}_{\mathrm{free}})) > 0$. Clearly, such n_0 exists, since $\lim_{n\to\infty} r_n = 0$ and $\mathcal{X}_{\mathrm{free}}$ has non-empty interior. Recall that ζ_d is the volume of the unit ball in the d-dimensional Euclidean space and that the connection radius of the PRM* algorithm is $r_n = \gamma_{\mathrm{PRM}}(\log n/n)^{1/d}$. Then, for all $n \geq n_0$

$$\mathbb{E}\big[M_n^{\mathrm{PRM}^*} \,\big|\, A\big] = \frac{\zeta_d\, \gamma_{\mathrm{PRM}}}{\mu(\mathrm{int}_{r_n}(\mathcal{X}_{\mathrm{free}}))} \log n.$$

On the other hand, given that $X_n \notin \mathrm{int}_{r_n}(\mathcal{X}_{\mathrm{free}})$, the r_n-ball centered at X_n intersects a fragment of $\mathcal{X}_{\mathrm{free}}$ that has volume less than the volume of an r_n-ball in the d-dimensional Euclidean space. Then, for all $n > n_0$, $\mathbb{E}\big[M_n^{\mathrm{PRM}^*} \,\big|\, A^c\big] \leq \mathbb{E}\big[M_n^{\mathrm{PRM}^*} \,\big|\, A\big]$.

Hence, for all $n \geq n_0$,

$$\mathbb{E}\left[\frac{M_n^{\mathrm{PRM}^*}}{\log n}\right] \leq \frac{\zeta_d\, \gamma_{\mathrm{PRM}}}{\mu(\mathrm{int}_{r_n}(\mathcal{X}_{\mathrm{free}}))} \leq \frac{\zeta_d\, \gamma_{\mathrm{PRM}}}{\mu(\mathrm{int}_{r_{n_0}}(\mathcal{X}_{\mathrm{free}}))}.$$

Next, consider the RRG. Recall that η is the parameter provided in the `Steer` procedure (see Section 3.1). Let D denote the diameter of the set $\mathcal{X}_{\mathrm{free}}$, i.e., $D := \sup_{x,x' \in \mathcal{X}_{\mathrm{free}}} \|x - x'\|$. Clearly, whenever $\eta \geq D$, $V^{\mathrm{PRM}^*} = V^{\mathrm{RRG}} = V^{\mathrm{RRT}^*}$ surely, and the claim holds.

To prove the claim when $\eta < D$, let C_n denote the event that for any point $x \in \mathcal{X}_{\mathrm{free}}$ the RRG algorithm has a vertex $x' \in V_n^{\mathrm{RRG}}$ such that $\|x - x'\| \leq \eta$. As shown in the proof of Theorem 36 (see Lemma 63), there exists $a, b > 0$ such that $\mathbb{P}(C_n^c) \leq a\, e^{-b\, n}$. Then,

$$\mathbb{E}\left[M_n^{\mathrm{RRG}}\right] = \mathbb{E}\left[M_n^{\mathrm{RRG}} \,\middle|\, C_n\right]\, \mathbb{P}(C_n) + \mathbb{E}\left[M_n^{\mathrm{RRG}} \,\middle|\, C_n^c\right]\, \mathbb{P}(C_n^c),$$

Clearly, $\mathbb{E}\left[M_n^{\mathrm{RRG}} \,\middle|\, C_n^c\right] \leq n$. Hence, the second term of the sum on the right hand side converges to zero as n approaches infinity. On the other hand, given that C_n holds, the new vertex that will be added to the graph at iteration n, if such a vertex is added at all, will be the same as the last sample, X_n. To complete the argument, given any set of n points placed inside $\mu(X_{\mathrm{free}})$, let N_n denote the number of points that are inside a ball of radius r_n that is centered at a point X_n sampled uniformly at random from $\mu(X_{\mathrm{free}})$. The expected number of points inside this ball is no more than $\frac{\zeta_d\, r_n^d}{\mu(X_{\mathrm{free}})}\, n$. Hence, $\mathbb{E}[M_n^{\mathrm{RRG}} \mid C_n] < \frac{\zeta_d\, \gamma_{\mathrm{PRM}}}{\mu(X_{\mathrm{free}})} \log n$, which implies the existence of a constant $\phi_1 \in \mathbb{R}_{\geq 0}$ such that $\limsup_{n\to\infty} \mathbb{E}[M_n^{\mathrm{RRG}}/(\log n)] \leq \phi_1$.

Finally, since $M_n^{\mathrm{RRT}^*} = M_n^{\mathrm{RRG}}$ holds surely, $\limsup_{n\to\infty} \mathbb{E}[M_n^{\mathrm{RRG}}/(\log n)] \leq \phi_1$ also. □

Trivially, $M_n^{k\text{-}\mathrm{PRM}^*} = M_n^{k\text{-}\mathrm{RRG}} = M_n^{k\text{-}\mathrm{RRT}^*} = k\, \log n$ for all n with $n/\log n > k$.

Complexity of the `CollisionFree` procedure In this section, complexity of the `CollisionFree` procedure in terms of the number of obstacles in the environment is analyzed, which is a widely-studied problem in the literature (see, e.g., Lin and Manocha (2004) for a survey). The main result is based on Six and Wood (1982), which shows that checking collision with m obstacles can be executed in $O(\log^d m)$ time using data structures based on spatial trees (see also Edelsbrunner and Maurer, 1981; Hopcroft et al., 1983).

Complexity of the `Nearest` procedure The nearest neighbor search problem has been widely studied in the literature, since it has many applications in, e.g., computer graphics, database systems, image processing, data mining, pattern recognition, etc. (Samet, 1989b,a). Clearly, a brute-force algorithm that examines every vertex runs in $O(n)$ time and requires $O(1)$ space. However, in many online real-time applications such as robotics, it is highly desirable to reduce the computation time of each iteration under sublinear bounds, e.g., in $O(\log n)$ time, especially for anytime algorithms that provide better solutions as the number of iterations increase.

Fortunately, existing algorithms for computing an "approximate" nearest neighbor, if not an exact one, are computationally very efficient. In the sequel, a vertex y is said to be an ε-approximate nearest neighbor of a point x if $\|y - x\| \leq (1 + \varepsilon)\, \|z - x\|$, where z is the true nearest neighbor of x. An approximate nearest neighbor can be computed using balanced-box decomposition (BBD) trees, which achieves $O(c_{d,\varepsilon} \log n)$ query time using $O(d\, n)$ space (Arya et al., 1999), where $c_{d,\varepsilon} \leq d\lceil 1{+}6d/\varepsilon\rceil^d$. This algorithm is computationally optimal in fixed dimensions, since it closely matches a lower bound for algorithms that use a tree structure stored in roughly linear space (Arya et al., 1999). Using approximate nearest neighbor computation in the context of both PRMs and RRTs was discussed very recently in Yershova and LaValle (2007); Plaku and Kavraki (2008).

Let $G = (V, E)$ be a graph with $V \subseteq \mathcal{X}$ and let $x \in \mathcal{X}$. The discussion above implies that the number of simple operations executed by the `Nearest`(G, x) procedure is $\Theta(\log |V|)$ in fixed dimensions, if the `Nearest` procedure is implemented using a tree structure that is stored in linear space.

Complexity of the `Near` procedure Problems similar to that solved by the `Near` procedure are also widely-studied in the literature, generally under the name of *range search problems*, as they have many applications in, for instance, computer graphics and spatial database systems (Samet, 1989a). In the worst case and in fixed dimensions, computing the exact set of vertices that reside in a ball of radius r_n centered at a query point x takes $O(n^{1-1/d} + m)$ time using k-d trees (Lee and Wong, 1977), where m is the number of vertices returned by the search (see also Chanzy et al. (2001) for an analysis of the average case).

Similar to the nearest neighbor search, computing approximate solutions to the range search problem is computationally easier. A range search algorithm is said to be ε-approximate if it returns all vertices that reside in the ball of size r_n and no vertices outside a ball of radius $(1+\varepsilon)\, r_n$, but may or may not return the vertices that lie outside the former ball and inside the latter ball. Computing ε-approximate solutions using BBD-trees requires $O(2^d \log n + d^2(3\sqrt{d}/\varepsilon)^{d-1})$ time when using $O(d\, n)$ space, in the worst case (Arya and Mount, 2000). Thus, in fixed dimensions, the complexity of this algorithm is $O(\log n + (1/\varepsilon)^{d-1})$, which is known to be optimal, closely matching a lower bound (Arya and Mount, 2000). More recently, algorithms that can provide trade-offs between time and space were also proposed (Arya et al., 2005).

Note that the `Near` procedure can be implemented as an approximate range search while maintaining the asymptotic optimality guarantee. Notice that the expected number of vertices returned by the `Near` procedure also does not change, except by a constant factor. Hence, the `Near` procedure can be implemented to run in order $\log n$ expected time in the limit and linear space in fixed dimensions.

Time complexity of the processing phase The following results characterize the asymptotic computational complexity of various sampling-based algorithms in terms of the number of simple operations such as comparisons, additions, and multiplications.

Let n denote the total number of iterations (or, alternatively, the number of samples), and m denote the number of obstacles in the environment. Then, by Lemmas 40 and 41, N_n^{PRM}, $N_n^{\mathrm{sPRM}} \in \Omega(n^2 \log^d m)$. In the k-nearest sPRM and RRT algorithms, $\Omega(\log n)$ time is spent on finding the (k-)nearest neighbor(s) and $\Omega(\log^d m)$ time is spent on collision checking at each iteration. Hence, $N_n^{k\text{-sPRM}}, N_n^{\mathrm{RRT}} \in \Omega(n \log n + n \log^d m)$.

In all the proposed algorithms, $O(\log n)$ time is spent on finding the near neighbors, and $\log n \log^d m$ time is spent on collision checking. Thus, $N_n^{\mathrm{ALG}} \in O(n\, \log n \log^d m)$ for $ALG \in \{\mathrm{PRM}^*, k\text{-PRM}^*, \mathrm{RRG}, k\text{-RRG}, \mathrm{RRT}^*, k\text{-RRT}^*\}$.

Time complexity of the query phase After algorithm ALG returns the graph G_n^{ALG}, the optimal path must be extracted from this graph using, e.g., Dijkstra's shortest path algorithm (Schrijver, 2003). In this section, the complexity of this operation, called the query phase, is discussed.

The following lemma yields the asymptotic computational complexity of computing shortest paths. Let $G = (V, E)$ be a graph. A length function $l : E \to \mathbb{R}_{>0}$ is a function that assigns each edge in E a positive length. Given a vertex $v \in V$, the shortest paths tree for G, l, and v is a graph $G' = (V, E')$, where $E' \subseteq E$ such that for any $v' \in V \setminus \{v\}$, there exists a unique path in G that starts from v and reaches v', moreover, this path is the optimal such path in G.

Lemma 43 (Complexity of shortest paths (Schrijver, 2003)) *Given a graph $G = (V, E)$, a length function $l : E \to \mathbb{R}_{>0}$, and a vertex $v \in V$, the shortest path tree for G, l, and v can be found in time $O(|V| \log(|V|) + |E|)$.*

It remains to determine the number of vertices and edges in $G_n^{\mathrm{ALG}} = (V_n^{\mathrm{ALG}}, E_n^{\mathrm{ALG}})$, for each algorithm ALG.

Trivially, $|E_n^{\mathrm{ALG}}| \in \Omega(n)$ holds for all the algorithms discussed in this paper, in particular, for $ALG \in \{\mathrm{PRM}, k\text{-sPRM}, \mathrm{RRT}\}$. For the sPRM algorithm, a stronger bound can be provided: $|E_n^{\mathrm{sPRM}}| \in \Omega(n^2)$. To prove this claim, consider the problem instance $(\mathcal{X}_{\mathrm{free}}, x_{\mathrm{init}}, \mathcal{X}_{\mathrm{goal}})$, where $\mathcal{X}_{\mathrm{free}} = \mathcal{X} = (0,1)^d$. Then, the straight path between any two vertices will be collision-free. Thus, the number of edges is exactly equal to the number of calls to the `CollisionFree` procedure. Then, the result follows from Lemma 41.

For the proposed algorithms, $|E_n^{\mathrm{PRM}^*}|, |E_n^{\mathrm{RRG}}| \in O(n \log n)$. Since the number of edges is always less than or equal to the total number of calls to the `CollisionFree` procedure, this claim follows directly from Lemma 42. Finally, $|E_n^{k\text{-}\mathrm{PRM}^*}|, |E_n^{k\text{-}\mathrm{RRG}}| \in O(n \log n)$ and $|E_n^{\mathrm{RRT}^*}|$, $|E_n^{k\text{-}\mathrm{RRT}^*}| \in O(n)$ all hold trivially.

Space complexity Space complexity of an algorithm ALG is defined as the amount of memory that is used by ALG to compute the graph $G_n^{\mathrm{ALG}} = (V_n^{\mathrm{ALG}}, E_n^{\mathrm{ALG}})$. Clearly, in all algorithms discussed in this paper, the space complexity is the size of G_n^{ALG}, i.e., $|V_n^{\mathrm{ALG}}| + |E_n^{\mathrm{ALG}}|$. Since the number of edges is at least as much as the number of vertices in G_n^{ALG} for all algorithms discussed in this paper, the space complexity of an algorithm, in this context, is the number edges in the graph that it returns, which was determined in the previous section.

5 Numerical Experiments

This section is devoted to an experimental study of the algorithms considered in the paper. All algorithms were implemented in C and run on a computer with 2.66 GHz processor and 4GB RAM running the Linux operating system. Unless otherwise noted, total variation of a path is its cost.

A first set of experiments were run to illustrate the different performance of k-nearest PRM and of PRM*. The k-nearest PRM and the PRM* algorithms were run alongside in two dimensional configuration-space and the cost of the best path in both algorithms is plotted versus the number of iterations in Figure 10. The k-nearest PRM does not converge to optimal solutions, unlike PRM*. The performance of the PRM* algorithm is also shown in configuration spaces of dimensions up to five in Figure 11.

The main bulk of the experiments were aimed at demonstrating the performance of the RRT* algorithm, especially in comparison with its “standard” counterpart, i.e., RRT. Three problem instances were considered. In the first two, the cost function is the Euclidean path length.

The first scenario includes no obstacles. Both algorithms are run in a square environment. The trees maintained by the algorithms are shown in Figure 12 at several stages. The figure illustrates that, in this case, the RRT algorithm does not improve the feasible solution to converge to an optimum solution. On the other hand, running the RRT* algorithm further improves the paths in the tree to lower cost ones. The convergence properties of the two algorithms are also investigated in Monte-Carlo runs. Both algorithms were run for 20,000 iterations 500 times and the cost of the best path in the trees were averaged for each iteration. The results are shown in Figure 13, which shows that in the limit the RRT algorithm has cost very close to a $\sqrt{2}$ factor the optimal solution (see LaValle and Kuffner (2009) for a similar result in a deterministic setting), whereas the RRT* converges to the optimal solution. Moreover, the variance over different RRT runs approaches

2.5, while that of the RRT* approaches zero. Hence, almost all RRT* runs have the property of convergence to an optimal solution, as expected.

In the second scenario, both algorithms are run in an environment in presence of obstacles. In Figure 14, the trees maintained by the algorithms are shown after 20,000 iterations. The tree maintained by the RRT* algorithm is also shown in Figure 15 in different stages. It can be observed that the RRT* first rapidly explores the state space just like the RRT. Moreover, as the number of samples increase, the RRT* improves its tree to include paths with smaller cost and eventually discovers a path in a different homotopy class, which reduces the cost of reaching the target considerably. Results of a Monte-Carlo study for this scenario is presented in Figure 16. Both algorithms were run alongside up until 20,000 iterations 500 times and cost of the best path in the trees were averaged for each iteration. The figures illustrate that all runs of the RRT* algorithm converges to the optimum, whereas the RRT algorithm is about 1.5 of the optimal solution on average. The high variance in solutions returned by the RRT algorithm stems from the fact that there are two different homotopy classes of paths that reach the goal. If the RRT luckily converges to a path of the homotopy class that contains an optimum solution, then the resulting path is relatively closer to the optimum than it is on average. If, on the other hand, the RRT first explores a path of the second homotopy class, which is often the case for this particular scenario, then the solution that RRT converges to is generally around twice the optimum.

Finally, in the third scenario, where no obstacles are present, the cost function is selected to be the line integral of a function, which evaluates to 2 in the high cost region, 1/2 in the low cost region, and 1 everywhere else. The tree maintained by the RRT* algorithm is shown after 20,000 iterations in Figure 17. Notice that the tree either avoids the high cost region or crosses it quickly, and vice-versa for the low-cost region. (Incidentally, this behavior corresponds to the well known Snell-Descartes law for refraction of light, see Rowe and Alexander (2000) for a path-planning application.)

To compare the running time, both algorithms were run alongside in an environment with no obstacles for up to one million iterations. Figure 18, shows the ratio of the running time of RRT* and that of RRT versus the number of iterations averaged over 50 runs. As expected from the complexity analysis of Section 4.3, this ratio converges to a constant value. A similar figure is produced for the second scenario and provided in Figure 19.

The RRT* algorithm was also run in a 5-dimensional state space. The number of iterations versus the cost of the best path averaged over 100 trials is shown in Figure 20. A comparison with the RRT algorithm is provided in the same figure. The ratio of the running times of the RRT* and the RRT algorithms is provided in Figure 21. The same experiment is carried out for a 10-dimensional configuration space. The results are shown in Figure 22.

6 Conclusion

This paper presented the results of a thorough analysis of sampling-based algorithms for optimal path planning. It is shown that broadly used algorithms from the literature, while probabilistically complete, are not asymptotically optimal, i.e., they will return a solution to the path planning problem with high probability if one exists, but the cost of the solution returned by the algorithm will not converge to the optimal cost as the number of samples increases. In particular, it is proven that the PRM and RRT algorithms are not asymptotically optimal. A simplified version of PRM is asymptotically optimal, but is computationally expensive. In addition, it is shown that certain heuristic versions of PRM are not only not asymptotically complete, but also not necessarily complete.

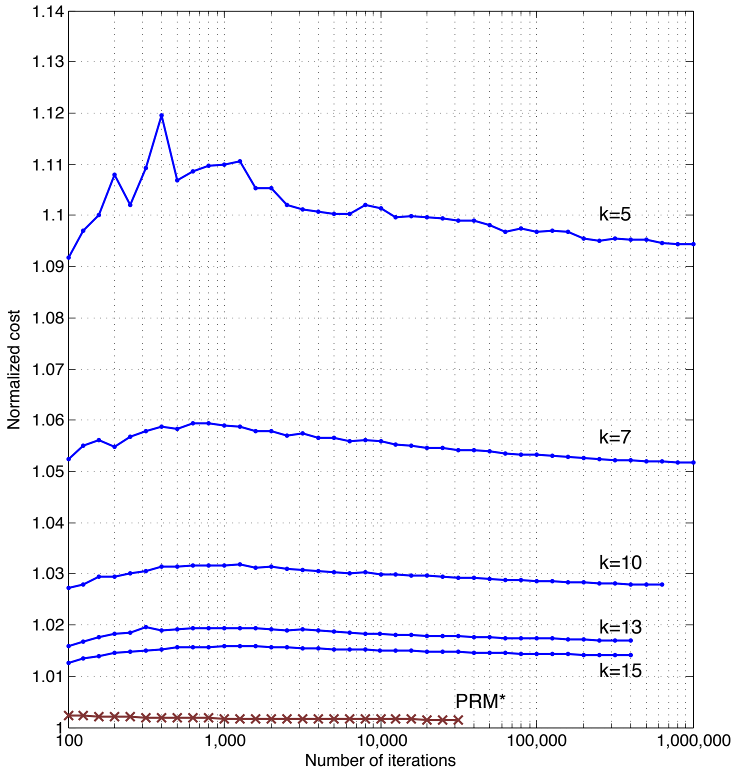

Figure 10: The cost of the best path in the k-nearest sPRM algorithm, and that in the PRM* algorithm are shown versus the number of iterations in simulation examples with no obstacles. The k-nearest sPRM algorithm was run for $k = 5, 7, 10, 13, 15$, each of which is shown separately in blue, and the PRM* algorithm is shown in red. The values are normalized so that the cost of the optimal path is equal to one. The iterations were stopped when the query phase of the algorithms exceeded the memory limit (approximately 4GB).

In order to address these limitations of existing algorithms, a number of new algorithms are introduced, and proven to be asymptotically optimal and computational efficient, with respect to probabilistically complete algorithms in this class. In other words, asymptotic optimality imposes only a constant factor increase in complexity with respect to probabilistic completeness. The first algorithm, called PRM*, is a variant of PRM, with a variable connection radius that scales as $\log(n)/n$, where n is the number of samples. In other words, the average number of connections made at each iteration is proportional to $\log(n)$. The second new algorithm, called RRG, incrementally builds a connected roadmap, augmenting the RRT algorithm with connections within a ball scaling as $\log(n)/n$. The third new algorithm, called RRT*, is a version of RRG that incrementally builds a tree. Experimental evidence that demonstrate the effectiveness of the algorithms proposed and support the theoretical claims were also provided.

A common theme in the paper is that, in order to ensure both asymptotic optimality and computational efficiency, connections between samples should be sought within balls of radius scaling as $\log(n)/n$. If these balls shrink faster as n increases, the algorithms are not asymptotically optimal (but may still be probabilistically complete); on the other hand, if these balls shrink slower, the complexity of the algorithms will suffer. On average, the proposed scaling laws will result in an average number of connections per iteration that is proportional to $\log(n)$. Hence, it is natural to

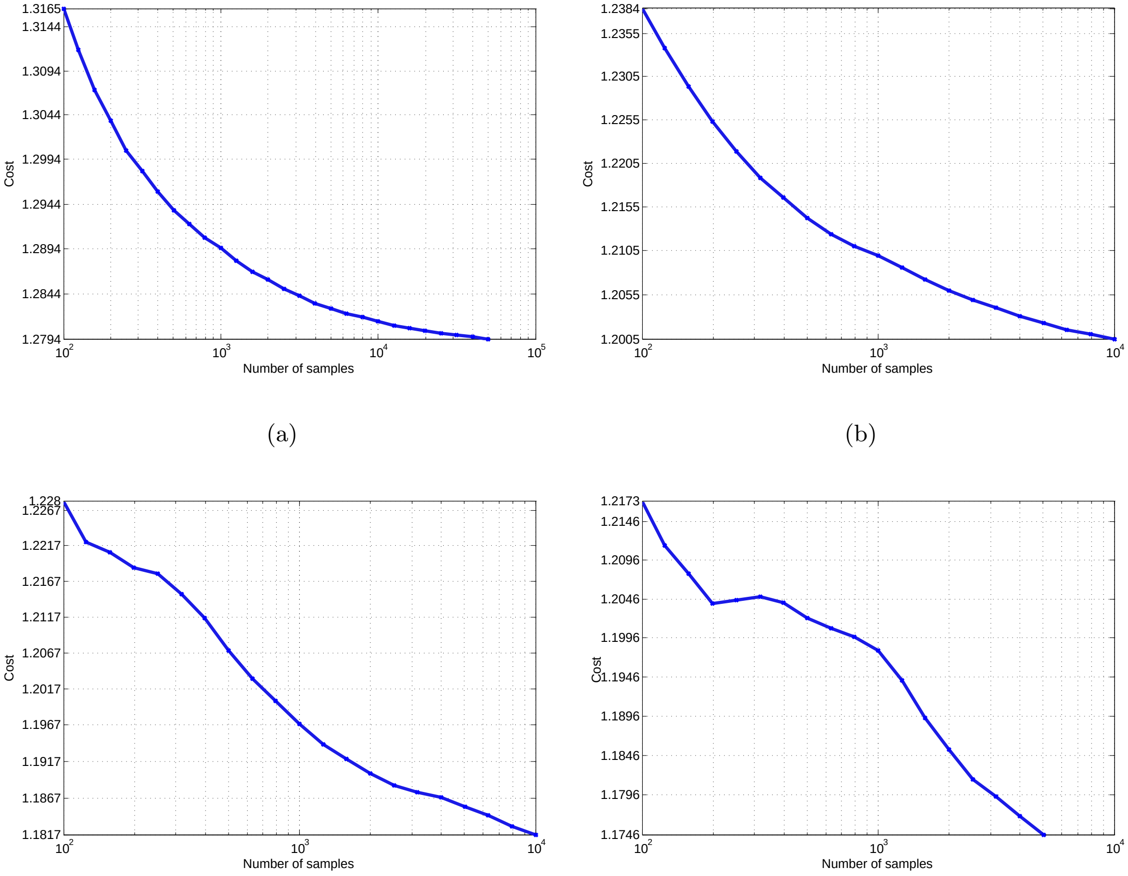

(a) (b)

(c) (d)

Figure 11: Cost of the best path in the PRM* algorithm is shown in up to 2, 3, 4, and 5 dimensional configuration spaces, in Figures (a), (b), (c), and (d), respectively. The initial condition and goal region are on opposite vertices of the unit cube $(0, 1)^d$. The obstacle region is a cube centered at $(0.5, 0.5, \ldots, 0.5)$ and has volume 0.5 in all cases.

consider variants of these algorithms that make connections to $k \log(n)$ neighbors surely. Indeed, it is shown that these algorithms do share the same asymptotic optimality and computational efficiency properties of their counterparts, as long as k is no smaller than a constant k^*_{RRG}. It is remarkable that this constant only depends on the dimension of the space, and is otherwise independent from the problem instance.

The analysis of the results in the paper relies on techniques used to analyze random geometric graphs. Indeed, the algorithms considered in this paper build graphs that have many characteristics in common with well known classes of random geometric graphs. Interestingly, such geometric graphs exhibit phase transition phenomena, including percolation and connectivity, for thresholds matching those found for probabilistic completeness and asymptotic optimality of sampling-based algorithms. This leads to a natural conjecture that a sampling-based path planning algorithm is probabilistically complete if and only if the underlying random geometric graph percolates, and is asymptotically optimal if and only if the underlying random geometric graph is connected.

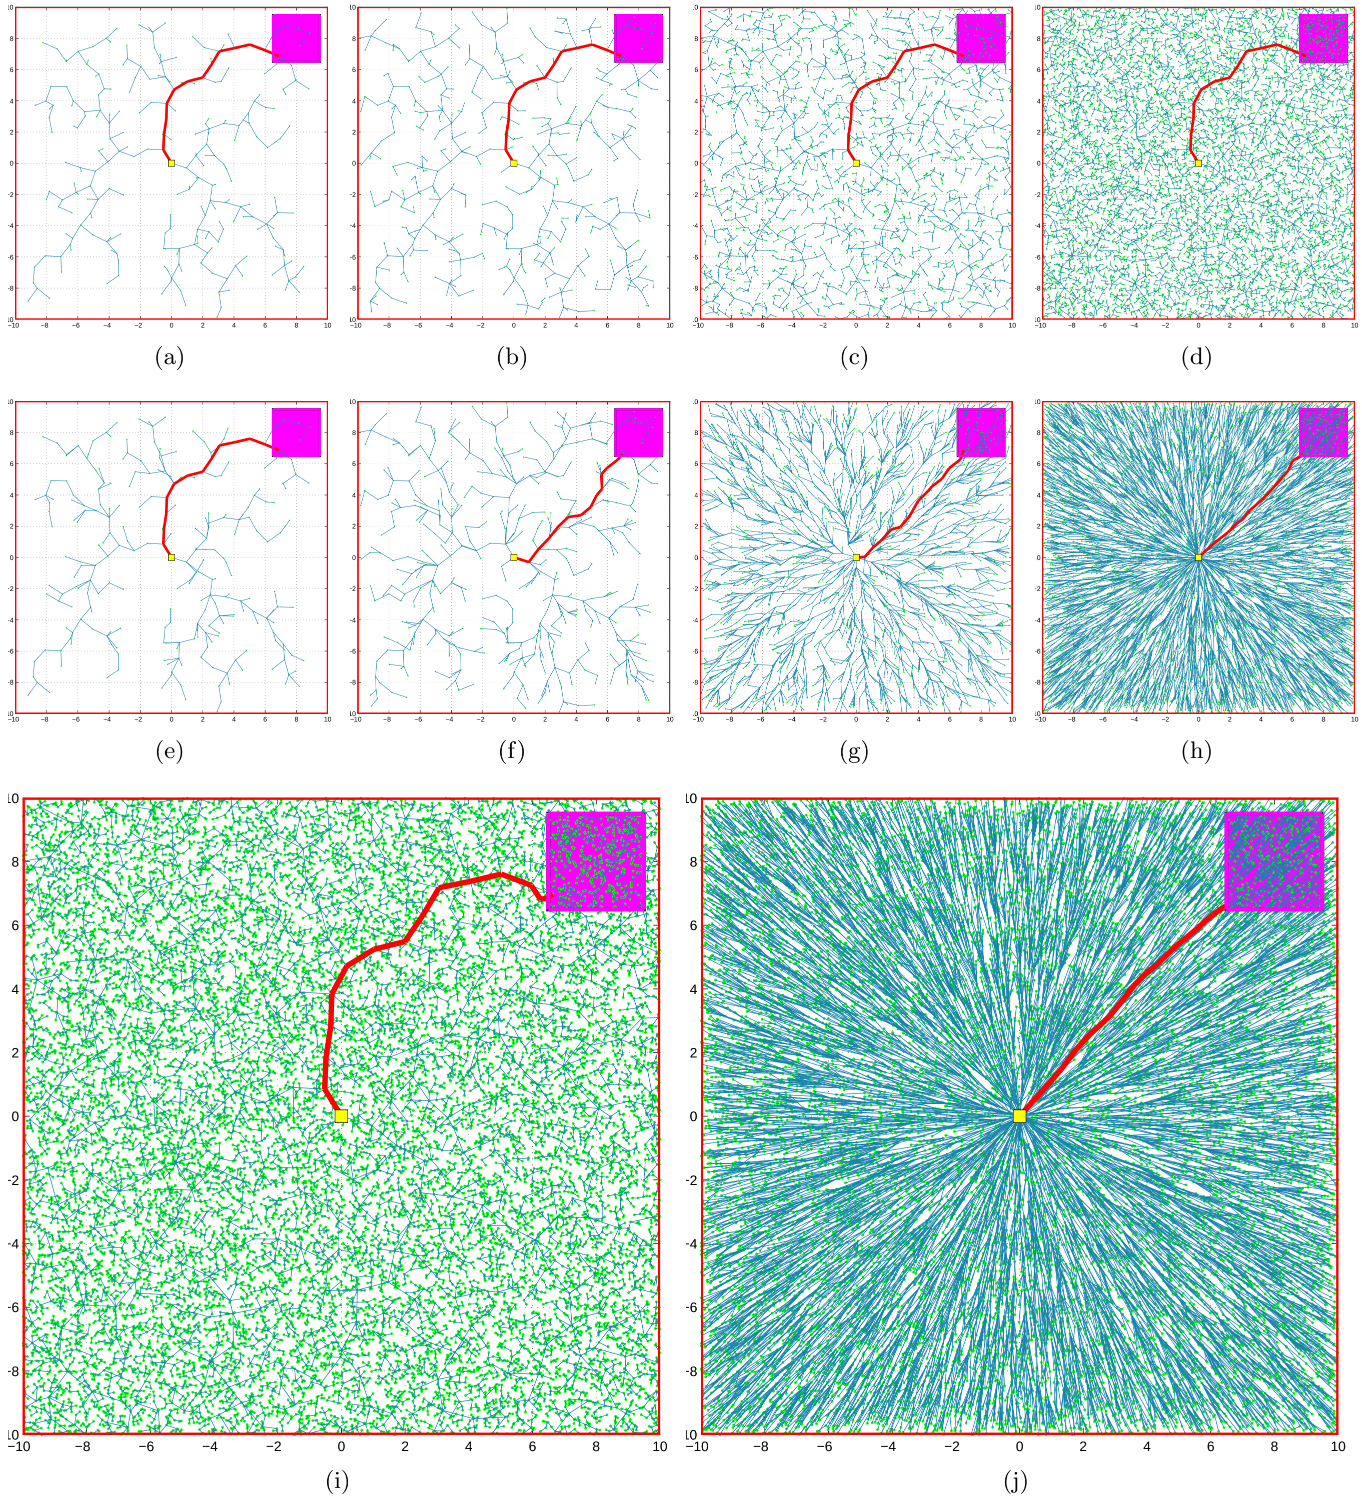

(a) (b) (c) (d)

(e) (f) (g) (h)

(i) (j)

Figure 12: A Comparison of the RRT* and RRT algorithms on a simulation example with no obstacles. Both algorithms were run with the same sample sequence. Consequently, in this case, the vertices of the trees at a given iteration number are the same for both of the algorithms; only the edges differ. The edges formed by the RRT algorithm are shown in (a)-(d) and (i), whereas those formed by the RRT* algorithm are shown in (e)-(h) and (j). The tree snapshots (a), (e) contain 250 vertices, (b), (f) 500 vertices, (c), (g) 2500 vertices, (d), (h) 10,000 vertices and (i), (j) 20,000 vertices. The goal regions are shown in magenta (in upper right). The best paths that reach the target in all the trees are highlighted with red.

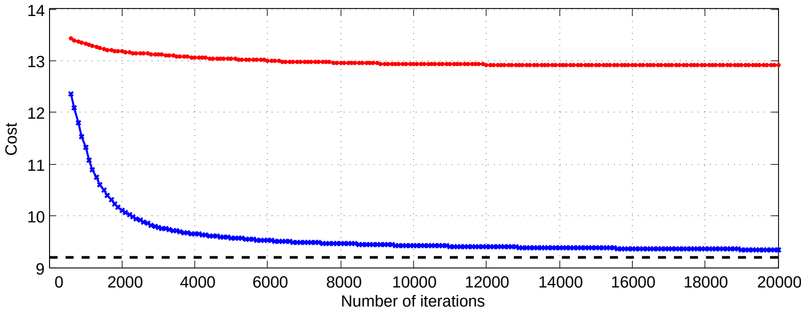

(a)

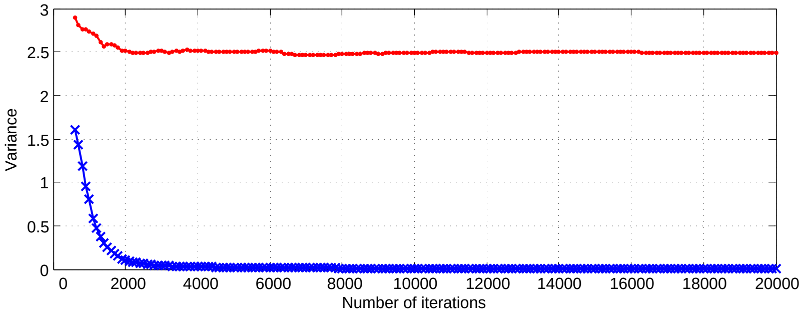

(b)

Figure 13: The cost of the best paths in the RRT (shown in red) and the RRT* (shown in blue) plotted against iterations averaged over 500 trials in (a). The optimal cost is shown in black. The variance of the trials is shown in (b).

The work presented in this paper can be extended in numerous directions. First of all, it would be of interest to establish broader connections between sampling-based path planning algorithms and random geometric graphs, e.g., by proving or disproving the conjecture above, and by possibly improving on current algorithms through a better understanding of the underlying mathematical objects. Similar analysis techniques can also be used to analyze other sampling-based path planning algorithms that were not analyzed in this paper, such as EST. In addition, it is of interest to investigate deterministic sampling-based algorithms, in which samples are generated using deterministic dense sequences of points with, e.g., low dispersion, as opposed to random sequences.

Second, it is of great practical interest to address motion planning problems subject to more complex constraints. For example, motion planning problems for mobile robots should consider the robot's dynamics, and hence differential constraints on the feasible trajectories (these are also called kino-dynamic planning problems). In addition, it is of interest to consider optimal planning problems in the presence of temporal/logic constraints on the trajectories, e.g., expressed using formal specification languages such as Linear Temporal Logic, or the μ-calculus. Such constraints correspond to, e.g., rules of the road constraints for autonomous ground vehicles, mission specifications for autonomous robots, and rules of engagement in military applications. Ultimately, incremental sampling-based algorithms with asymptotic optimality properties may provide the ba-

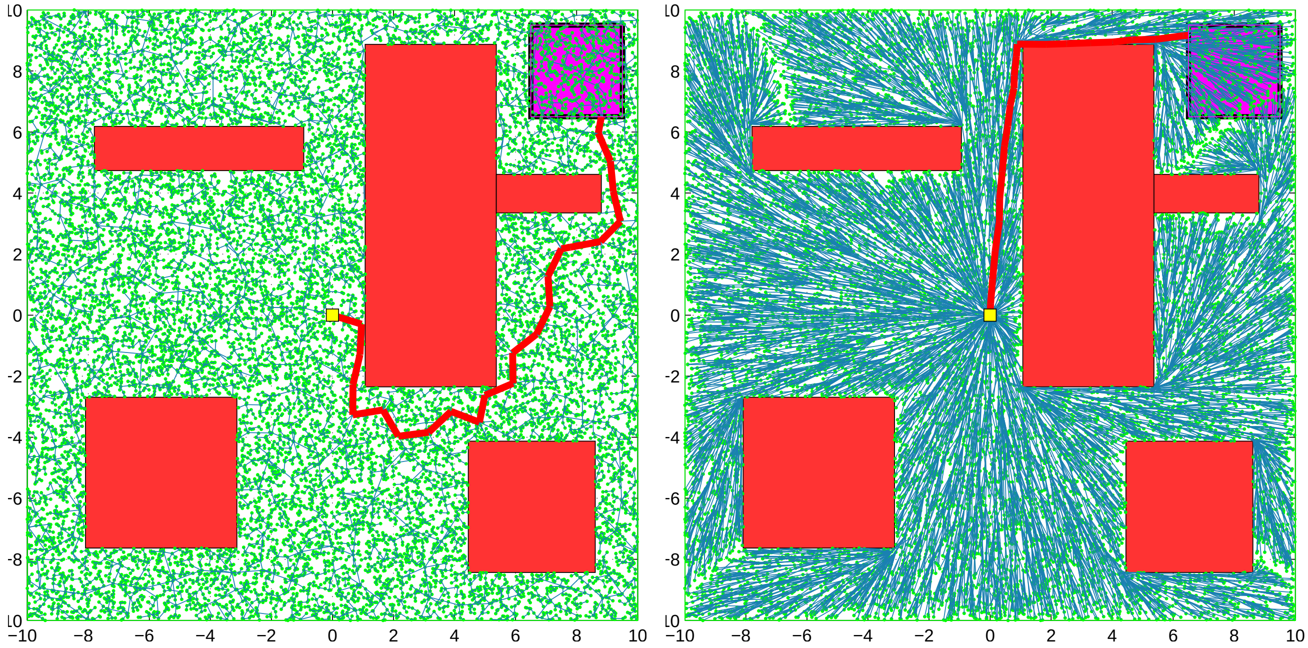

(a)

(b)

Figure 14: A Comparison of the RRT (shown in (a)) and RRT* (shown in (b)) algorithms on a simulation example with obstacles. Both algorithms were run with the same sample sequence for 20,000 samples. The cost of best path in the RRT and the RRG were 21.02 and 14.51, respectively.

sic elements for the on-line solution of differential games, as those arising when planning in the presence of dynamic obstacles.

Finally, it is noted that the proposed algorithms may have applications outside of the robotic motion planning domain. In fact, the class of sampling-based algorithm described in this paper can be readily extended to deal with problems described by partial differential equations, such as the eikonal equation and the Hamilton-Jacobi-Bellman equation.

Acknowledgments

The authors are grateful to Professors M.S. Branicky, G.J. Gordon, and S. LaValle, as well as the anonymous reviewers, for their insightful comments on draft versions of this paper. This research was supported in part by the Michigan/AFRL Collaborative Center on Control Sciences, AFOSR grant #FA 8650-07-2-3744, and by the National Science Foundation, grant CNS-1016213.

References

M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions*. Dover, 1964.

R. Alterovitz, S. Patil, and A. Derbakova. Rapidly-exploring roadmaps: Weighing exploration vs. reginement in optimal motion planning. In *IEEE Conference on Robotics and Automation (ICRA)*, 2011.

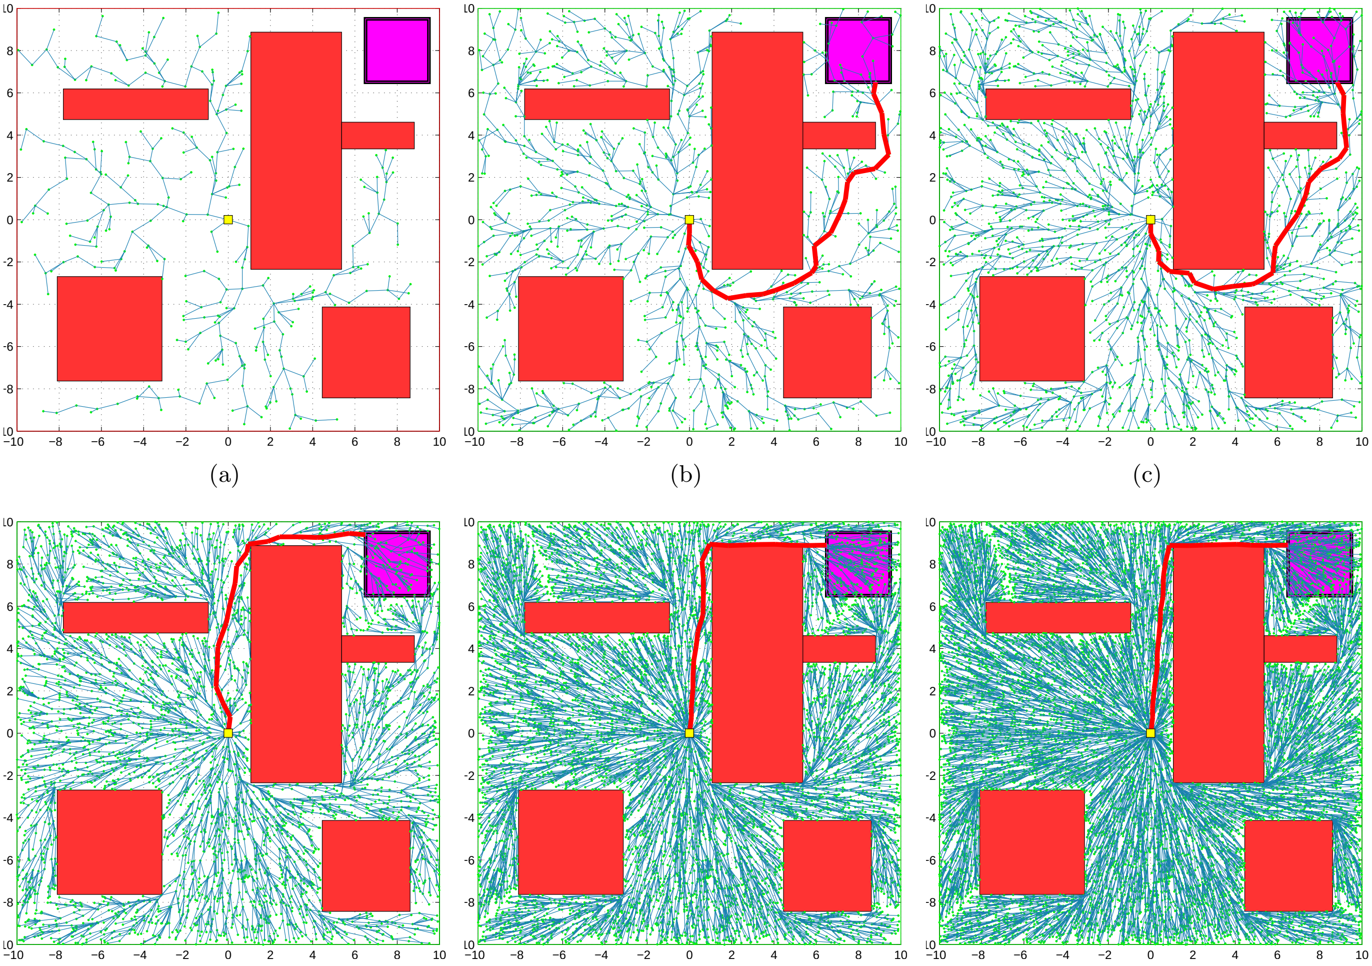

(a) (b) (c)

(d) (e) (f)

Figure 15: RRT* algorithm shown after 500 (a), 1,500 (b), 2,500 (c), 5,000 (d), 10,000 (e), 15,000 (f) iterations.

S. Arya and D. M. Mount. Approximate range searching. *Computational Geometry: Theory and Applications*, 17:135–163, 2000.

S. Arya, D. M. Mount, R. Silverman, and A. Y. Wu. An optimal algorithm for approximate nearest neighbor search in fixed dimensions. *Journal of the ACM*, 45(6):891–923, November 1999.

S. Arya, T. Malamatos, and D. M. Mount. Space-time tradeoffs for approximate spherical range counting. In *Symposium on Discrete Algorithms*, 2005.

P. Balister, B. Bollobás, A. Sarkar, and M. Walters. Connectivity of random k-nearest neighbour graphs. *Advances in Applied Probability*, 37:1–24, 2005.

P. Balister, B. Bollobás, and A. Sarkar. Percolation, connectivity, coverage and colouring of random geometric graphs. In B. Bollobás, R. Kozma, and D. Miklós, editors, *Handbook of Large-Scale Random Networks*, volume 18 of *Bolyai Society Mathematical Studies*, chapter 2, pages 117–142. Springer, 2009a.

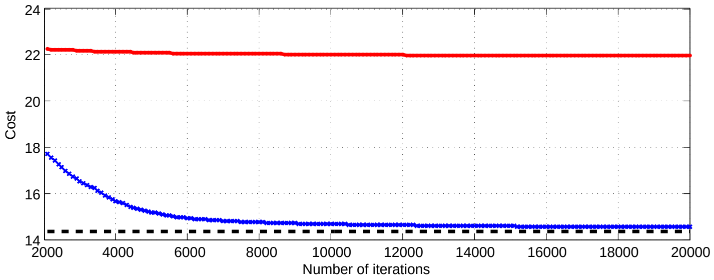

(a)

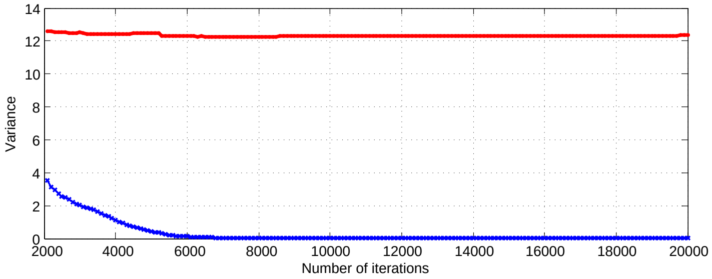

(b)

Figure 16: An environment cluttered with obstacles is considered. The cost of the best paths in the RRT (shown in red) and the RRT* (shown in blue) plotted against iterations averaged over 500 trials in (a). The optimal cost is shown in black. The variance of the trials is shown in (b).

P. Balister, B. Bollobás, A. Sarkar, and M. Walters. A critical constant for the k nearest-neighbour model. *Advances in Applied Probability*, 41(1):1–12, 2009b.

J. Barraquand and J. C. Latombe. Robot motion planning: A distributed representation approach. *International Journal of Robotics Research*, 10(6):628–649, 1993.

J. Barraquand, L. E. Kavraki, J. C. Latombe, T. Li, R. Motwani, and P. Raghavan. A random sampling scheme for path planning. *International Journal of Robotics Research*, 16:759–774, 1997.

D. Berenson, J. Kuffner, and H. Choset. An optimization approach to planning for mobile manipulation. In *IEEE International Conference on Robotics and Automation*, 2008.

D. Berenson, T. Simeon, and S. Srinivasa. Addressing cost-space chasms in manipulation planning. In *IEEE Conference on Robotics and Automation (ICRA)*, 2011.

A. Bhatia and E. Frazzoli. Incremental search methods for reachability analysis of continuous and hybrid systems. In R. Alur and G.J. Pappas, editors, *Hybrid Systems: Computation and Control*, number 2993 in Lecture Notes in Computer Science, pages 142–156. Springer-Verlag, Philadelphia, PA, March 2004.

B. Bollobás. *Random Graphs*. Cambridge University Press, second edition, 2001.

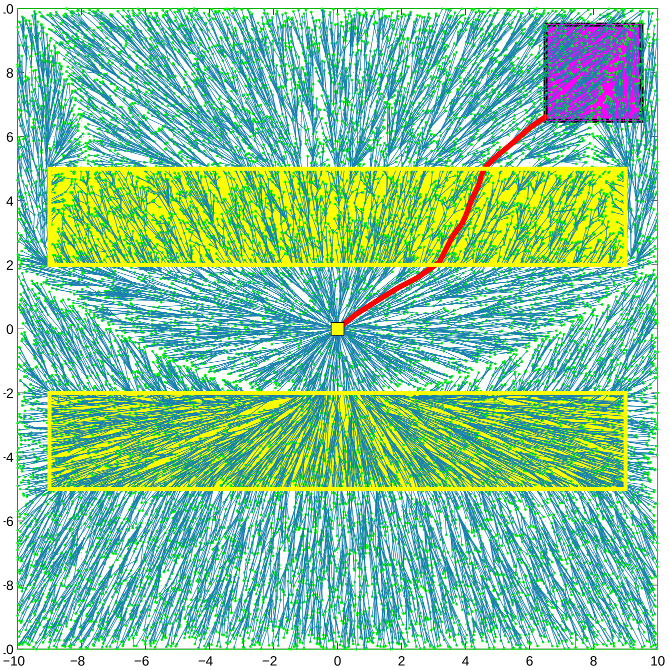

Figure 17: RRT* algorithm at the end of iteration 20,000 in an environment with no obstacles. The upper yellow region is the high-cost region, whereas the lower yellow region is low-cost.

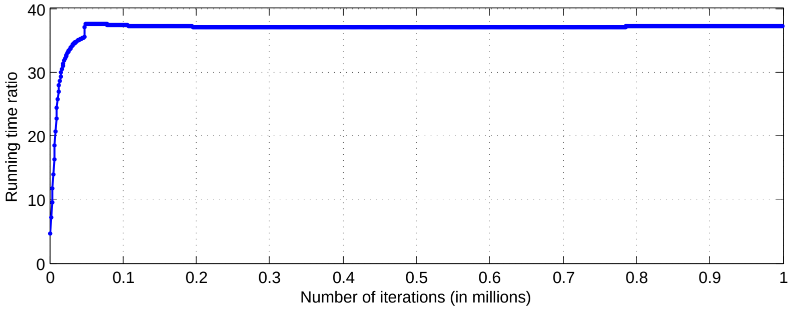

Figure 18: A comparison of the running time of the RRT* and the RRT algorithms. The ratio of the running time of the RRT* over that of the RRT up until each iteration is plotted versus the number of iterations.

B. Bollobás and O. M. Riordan. *Percolation.* Cambridge University Press, 2006.

M. S. Branicky, S. M. LaValle, K. Olson, and L. Yang. Quasi-randomized path planning. In *IEEE Conference on Robotics and Automation*, 2001.

M. S. Branicky, M. M. Curtis, J. A. Levine, and S. B. Morgan. RRTs for nonlinear, discrete, and hybrid planning and control. In *IEEE Conference on Decision and Control*, 2003.

M. S. Branicky, M. M. Curtis, J. Levine, and S. Morgan. Sampling-based planning, control, and

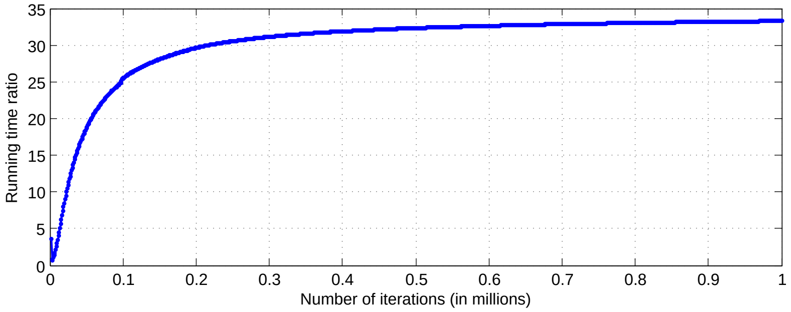

Figure 19: A comparison of the running time of the RRT* and the RRT algorithms in an environment with obstacles. The ratio of the running time of the RRT* over that of the RRT up until each iteration is plotted versus the number of iterations.

verification of hybrid systems. *IEEE Proc. Control Theory and Applications*, 153(5):575–590, Sept. 2006.

R. Brooks and T. Lozano-Perez. A subdivision algorithm in configuration space for findpath with rotation. In *International Joint Conference on Artificial Intelligence*, 1983.

J. Bruce and M.M. Veloso. *Real-Time Randomized Path Planning for Robot Navigation*, volume 2752 of *Lecture Notes in Computer Science*, chapter RoboCup 2002: Robot Soccer World Cup VI, pages 288–295. Springer, 2003.

A. Bry and N. Roy. Rapidly-exploring random belief trees for motion planning under uncertainty. In *IEEE Conference on Robotics and Automation (ICRA)*, 2011.

J. Canny. *The Complexity of Robot Motion Planning.* MIT Press, 1988.

J. Canny and J. H. Reif. New lower bound techniques for robot motion planning problems. In *IEEE Symposium on Foundations of Computer Science (FoCS)*, pages 49–60, Los Angeles, CA, 1987.

P. Chanzy, L. Devroye, and C. Zamora-Cura. Analysis of range search for random k-d trees. *Acta Informatica*, 37:355–383, 2001.

H. Choset, K.M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L.E. Kavraki, and S. Thrun. *Principles of Robot Motion: Theory, Algorithms, and Implementations.* MIT Press, Boston, MA, 2005.

J. Cortes, L. Jailet, and T. Simeon. Molecular disassembly with RRT-like algorithms. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2007.

D. Dolgov, S. Thrun, M. Montemerlo, and J. Diebel. *Experimental Robotics*, chapter Path Planning for Autonomous Driving in Unknown Environments, pages 55–64. Springer, 2009.

D. P. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms.* Cambridge University Press, 2009.

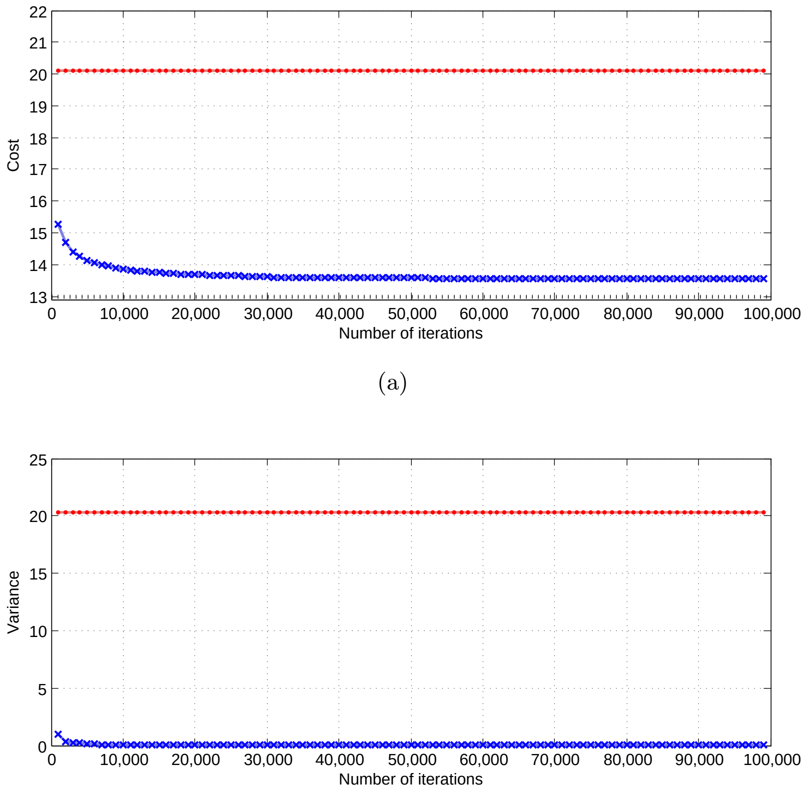

(a)

(b)

Figure 20: The cost of the best paths in the RRT (shown in red) and the RRT* (shown in blue) run in a 5 dimensional obstacle-free configuration space plotted against iterations averaged over 100 trials in (a). The optimal cost is shown in black. The variance of the trials is shown in (b).

H. Edelsbrunner and H. A. Maurer. On the intersection of orthogonal objects. *Information Processing Letters*, 13(4,5):177–181, April 1981.

D Eppstein, MS Paterson, and F F Yao. On nearest-neighbor graphs. *Discrete and Computational Geometry*, 17:263–282, Jan 1997. URL `http://www.springerlink.com/index/RM3FJ00T9AD4WBX9.pdf`.

D. Ferguson and A. Stentz. Anytime RRTs. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2006.

P. W. Finn and L. E. Kavraki. Computational approaches to drug design. *Algorithmica*, 25:347–371, 1999.

E. Frazzoli, M. A. Dahleh, and E. Feron. Real-time motion planning for agile autonomous vehicles. *Journal of Guidance, Control, and Dynamics*, 25(1):116–129, 2002.

S. S. Ge and Y.J. Cui. Dynamic motion planning for mobile robots using potential field method. *Autonomous Robots*, 13(3):207–222, 2002.

E. N. Gilbert. Random plane networks. *J. Soc. Indust. Appl. Math.*, 9(4):533–543, 1961.

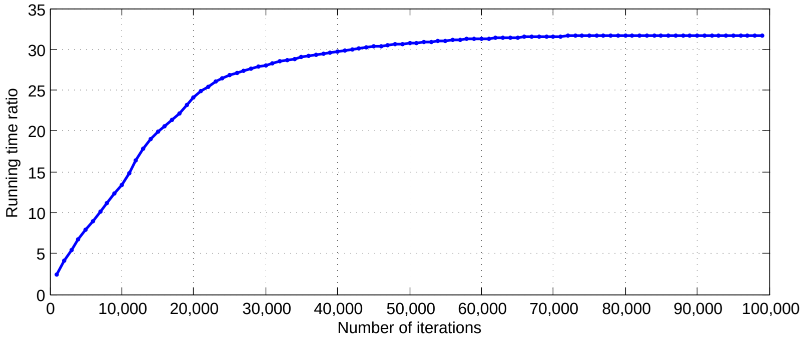

(a)

Figure 21: The ratio of the running time of the RRT and the RRT* algorithms is shown versus the number of iterations.

G. Grimmett and D. Stirzaker. *Probability and Random Processes.* Oxford University Press, Third edition, 2001.

P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity in wireless networks. In W. M. McEneany, G. Yin, and Q. Zhang, editors, *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, pages 547–566. Birkhäuser, Boston, MA, 1998.

P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. on Information Theory*, 46:388–404, 2000.

N. Henze. On the fraction of points with specified nearest-neighbour interrelations and degree of attraction. *Advances in Applied Probability*, 19:873–895, 1987.

J.E. Hopcroft, J.T. Schwartz, and M. Sharir. Efficient detection of intersections among spheres. *Int. Journal of Robotics Research*, 2:77–80, 1983.

D. Hsu, J. C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. In *IEEE Conference on Robotics and Automation*, 1997.

D. Hsu, J. C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. *Int. J. of Computational Geometry and Applications*, 9(4&5):495–512, 1999.

D. Hsu, R. Kindel, J. C. Latombe, and S. Rock. Randomized kinodynamic motion planning with moving obstacles. *International Journal of Robotics Research*, 21(3):233–255, 2002.

D. Hsu, J. C. Latombe, and H. Kurniawati. On the probabilistic foundations of probabilistic roadmap planning. *International Journal of Robotics Research*, 25:7, 2006.

L. Jaillet, J. Cortes, and T .Simeon. Sampling-based path planning on configuration-space costmaps. *IEEE Transactions on Robotics*, 26(4):635–646, August 2010.

S. Karaman and E. Frazzoli. Optimal kinodynamic motion planning using incremental sampling-based methods. In *IEEE Conf. on Decision and Control*, 2010a.

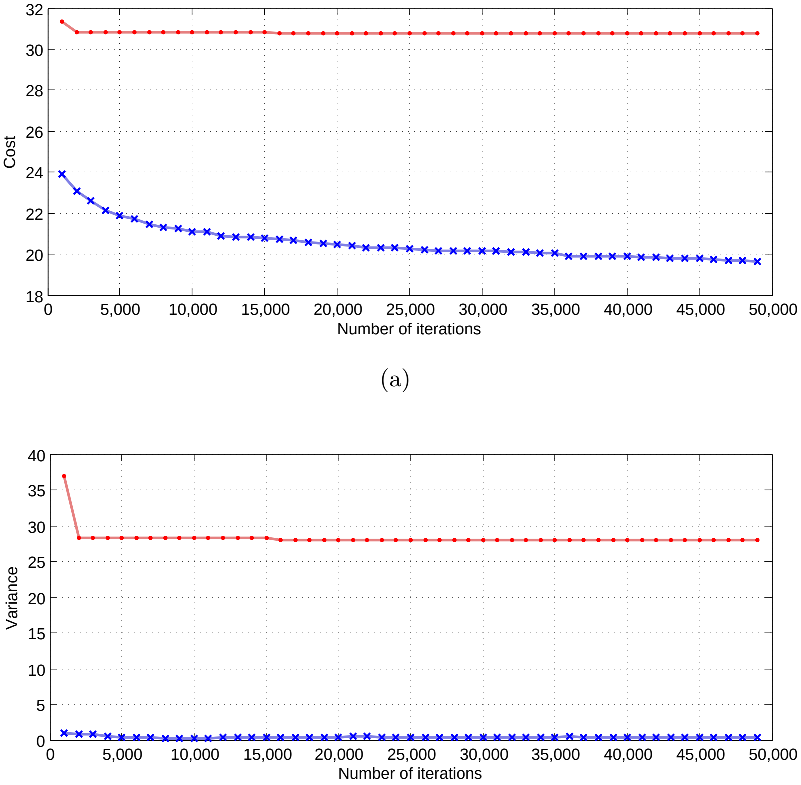

(a)

(b)

Figure 22: The cost of the best paths in the RRT (shown in red) and the RRT* (shown in blue) run in a 10 dimensional configuration space involving obstacles plotted against iterations averaged over 25 trials in (a). The variance of the trials is shown in (b).

S. Karaman and E. Frazzoli. Incremental sampling-based algorithms for optimal motion planning. In *Robotics: Science and Systems (RSS)*, 2010b.

S. Karaman and E. Frazzoli. Incremental sampling-based algorithms for a class of pursuit-evasion games. In *Workshop on Algorithmic Foundations of Robotics (WAFR)*, pages 71–87, 2010c.

S. Karaman, M. Walter, A. Perez, E. Frazzoli, and S. Teller. Anytime motion planning using the RRT*. In *IEEE Conference on Robotics and Automation (ICRA)*, 2011.

L. Kavraki and J. C. Latombe. Randomized preprocessing of configuration space for fast path planning. In *IEEE International Conference on Robotics and Automation*, 1994.

L. E. Kavraki, P. Svestka, J. C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.

L. E. Kavraki, M. N. Kolountzakis, and J. C. Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Transactions on Roborics and Automation*, 14(1):166–171, 1998.

O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *International Journal of Robotics Research*, 5(1):90–98, 1986.

Y. Koren and J. Borenstein. Potential field methods and their inherent limitations for mobile robot navigation. In *IEEE Conference on Robotics and Automation*, 1991.

E. Koyuncu, N.K. Ure, and G. Inalhan. Integration of path/manuever planning in complex environments for agile maneuvering UCAVs. *Jounal of Intelligent and Robotic Systems*, 57(1–4): 143–170, 2010.

J. J. Kuffner and S. M. LaValle. RRT-connect: An efficient approach to single-quert path planning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2000.

J. J. Kuffner, S. Kagami, K. Nishiwaki, M. Inaba, and H. Inoue. Dynamically-stable motion planning for humanoid robots. *Autonomous Robots*, 15:105–118, 2002.

Y. Kuwata, J. Teo, G. Fiore, S. Karaman, E. Frazzoli, and J.P. How. Real-time motion planning with applications to autonomous urban driving. *IEEE Transactions on Control Systems*, 17(5): 1105–1118, 2009.

A. L. Ladd and L. Kavraki. Measure theoretic analysis of probabilistic path planning. *IEEE Transactions on Robotics and Automation*, 20(2):229–242, 2004.

J. C. Latombe. *Robot Motion Planning.* Kluwer Academic Publishers, Boston, MA, 1991.

J. C. Latombe. Motion planning: A journey of robots, molecules, digital actors, and other artifacts. *International Journal of Robotics Research*, 18(11):1119–1128, 1999.

S. M. LaValle. *Planning Algorithms.* Cambridge University Press, 2006.

S. M. LaValle and J. J. Kuffner. Randomized kinodynamic planning. *International Journal of Robotics Research*, 20(5):378–400, May 2001.

S. M. LaValle and J. J. Kuffner. Space filling trees. Technical Report CMU-RI-TR-09-47, Carnegie Mellon University, The Robotics Institute, 2009.

S. M. LaValle, M. S. Branicky, and S. R. Lindemann. On the relationship between classical grid search and probabilistic roadmaps. *International Journal of Robotics Research*, 23(7–8):673–692, 2004.

D. T. Lee and C. K. Wong. Worst-case analysis for region and partial region searches in multidimensional binary search trees and quad trees. *Acta Informatica*, 9:23–29, 1977.

M. Likhachev and D. Ferguson. Planning long dynamically-feasible maneuvers for autonomous vehicles. *International Journal of Robotics Research*, 28(8):933–945, 2009.

M. Likhachev, G. Gordon, and S. Thrun. Anytime A* with provable bounds on sub-optimality. In *Advances in Neural Information Processing Systems*, 2004.

M. Likhachev, D. Ferguson, G. Gordon, A. Stentz, and S. Thrun. Anytime search in dynamic graphs. *Artificial intelligence Journal*, 172(14):1613–1643, 2008.

M. C. Lin and D. Manocha. Collision and proximity queries. In J.E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry.* Chapman and Hall/CRC, second edition, 2004.

S. R. Lindemann and S. M. LaValle. Current issues in sampling-based motion planning. In P. Dario and R. Chatila, editors, *Eleventh International Symposium on Robotics Research*, pages 36–54. Springer, 2005.

Y. Liu and N.I. Badler. Real-time reach planning for animated characters using hardware acceleration. In *IEEE International Conference on Computer Animation and Social Characters*, pages 86–93, 2003.

T. Lozano-Perez and M. A. Wesley. An algorithm for planning collision-free paths among polyhedral obstacles. *Communications of the ACM*, 22(10):560–570, 1979.

B. Luders, S. Karaman, E. Frazzoli, and J. P. How. Bounds on tracking error using closed-loop rapidly-exploring random trees. In *American Control Conference*, 2010.

R. Meester and R. Roy. *Continuum Percolation.* Cambridge University Press, 1996.

J. R. Munkres. *Topology.* Prentice Hall, second edition, 2000.

O. Nechushtan, B. Raveh, and D. Halperin. Sampling-diagram automata: a tool for analyzing path quality in tree planners. In D. Hsu, V. Isler, J. C. Latombe, and M.C. Lin, editors, *Algorithmic Foundations of Robotics IX*, volume 68 of *Springer tracts in advanced robotics*, pages 285–301. Springer, 2010.

H. Niederreiter. *Random Number Generation and Quasi-Monte-Carlo Methods.* Society for Industrial and Applied Mathematics, 1992.

M. D. Penrose. *Random Geometric Graphs.* Oxford University Press, 2003.

E. Plaku and L. E. Kavraki. Quantitative analysis of nearest-neighbors search in high-dimensional sampling-based motion planning. In *Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2008.

E. Plaku, K.E. Bekris, B.Y. Chen, A.M. Ladd, and L.E. Kavraki. Sampling-based roadmap of trees for parallel motion planning. *IEEE Transactions on Robotics*, 21:597–608, 2005.

S. Prentice and N. Roy. The belief roadmap: Efficient planning in blief space by factoring the covariance. *International Journal of Robotics Research*, 28(11–12):1448–1465, 2009.

J. Quintanilla, S. Torquato, and R.M. Ziff. Efficient measurement of the percolation threshold for fully penetrable discs. *Journal of Physics A*, 33(42):L399–L407, 2000.

J. H. Reif. Complexity of the mover's problem and generalizations. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, 1979.

S. I. Resnick. *A probability path.* Birkhäuser, 1999.

E. Rimon and D. E. Koditschek. Exact robot navigation using artificial potential fields. *IEEE Transactions on Robotics and Automation*, 8(5):501–518, 1992.

N. C. Rowe and R. S. Alexander. Finding optimal-path maps for path planning across weighted regions. *The International Journal of Robotics Research*, 19:83–95, 2000.

M. Sahimi. *Applications of Percolation Theory.* Taylor & Francis, 1994.

H. Samet. *Design and Analysis of Spatial Data Structures.* Addison-Wesley, 1989a.

H. Samet. *Applications of Spatial Data Structures: Computer Graphics, Image Processesing and Gis.* Addison-Wesley, 1989b.

A. Schrijver. *Combinatorial Optimization*, volume A. Springer, 2003.

J. T. Schwartz and M. Sharir. On the 'piano movers' problem: II. general techniques for computing topological properties of real algebraic manifolds. *Advances in Applied Mathematics*, 4:298–351, 1983.

A. Shkolnik, M. Levashov, I. R. Manchester, and R. Tedrake. Bounding on rough terrain with the LittleDog robot. Submitted for publication, 2011.

H. Six and D. Wood. Counting and reporting intersections of D-ranges. *IEEE Trans. on Computers*, pages 46–55, 1982.

D. Stentz. The focussed D* algorithm for real-time replanning. In *International Joint Conference on Artificial Intelligence*, 1995.

M. Stilman, J. Schamburek, J. Kuffner, and T. Asfour. Manipulation planning among movable obstacles. In *IEEE International Conference on Robotics and Automation*, 2007.

D. Stoyan, W. S. Kendall, and J. Mecke. *Stochastic Geometry and Its Applications.* John Wiley & Sons, 1995.

R. Tedrake, I. R. Manchester, M. M. Tobekin, and J. W. Roberts. LQR-trees: Feedback motion planning via sums of squares verification. *International Journal of Robotics Research (to appear)*, 2010.

S. Teller, M. R. Walter, M. Antone, A. Correa, R. Davis, L. Fletcher, E. Frazzoli, J. Glass, J.P. How, A. S. Huang, J. Jeon, S. Karaman, B. Luders, N. Roy, and T. Sainath. A voice-commandable robotic forklift working alongside humans in minimally-prepared outdoor environments. In *IEEE International Conference on Robotics and Automation*, 2010.

C. Urmson and R. Simmons. Approaches for heuristically biasing RRT growth. In *Proceedings of the IEEE/RSJ International Conference on Robotics and Systems (IROS)*, 2003.

A. R. Wade. Explicit laws of large numbers for random nearest-neighbor-type graphs. *Advances in Applied Probability*, 39:326–342, 2007.

A. R. Wade. Asymptotic theory for the multidimensional random on-line nearest-neighbour graph. *Stochastic Processes and their Applications*, 119(6):1889–1911, 2009.

N. A. Wedge and M.S. Branicky. On heavy-tailed runtimes and restarts in rapidly-exploring random trees. In *Twenty-third AAAI Conference on Artificial Intelligence*, 2008.

F. Xue and P. R. Kumar. The number of neighbors needed for connectivity of wireless networks. *Wireless Networks*, 10:169–181, 2004.

A. Yershova and S. M. LaValle. Improving motion-planning algorithms by efficient nearest-neighbor searching. *IEEE Transactions on Robotics*, 23(1):151–157, 2007.

A. Yershova and S. M. LaValle. Motion planning in highly constrained spaces. Technical report, University of Illinois at Urbana-Champaign, 2008.

M. Zucker, J. J. Kuffner, and M. S. Branicky. Multiple RRTs for rapid replanning in dynamic environments. In *IEEE Conference on Robotics and Automation*, 2007.

Appendix

A Notation

Let $\mathbb{N}$ denote the set of positive integers and $\mathbb{R}$ denote the set of reals. Let $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, and $\mathbb{R}_{>0}$, $\mathbb{R}_{\geq 0}$ denote the sets of positive and non-negative reals, respectively. A sequence on a set A is a mapping from $\mathbb{N}$ to A, denoted as $\{a_i\}_{i\in\mathbb{N}}$, where $a_i \in A$ is the element that $i \in \mathbb{N}$ is mapped to. Given $a, b \in \mathbb{R}$, closed and open intervals between a and b are denoted by $[a, b]$ and (a, b), respectively. The Euclidean norm is denoted by $\|\cdot\|$. Given a set $\mathcal{X} \subset \mathbb{R}^d$, the closure of $\mathcal{X}$ is denoted by $\mathrm{cl}(\mathcal{X})$. The closed ball of radius $r > 0$ centered at $x \in \mathbb{R}^d$, i.e., , i.e., $\{y \in \mathbb{R}^d \mid \|y - x\| \leq r\}$, is denoted as $\mathcal{B}_{x,r}$; $\mathcal{B}_{x,r}$ is also called the r-ball centered at x. Given a set $\mathcal{X} \subseteq \mathbb{R}^d$, the Lebesgue measure of X is denoted by $\mu(\mathcal{X})$. The Lebesgue measure of a set is also referred to as its volume. The volume of the unit ball in $\mathbb{R}^d$, is denoted by ζ_d, i.e., $\zeta_d = \mu(\mathcal{B}_{0,1})$. The letter e is used to denote the base of the natural logarithm, also called Euler's number.

Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a sample space, $\mathcal{F} \subseteq 2^\Omega$ is a $\sigma-$algebra, and $\mathbb{P}$ is a probability measure, an event A is an element of $\mathcal{F}$. The complement of an event A is denoted by A^c. Given a sequence of events $\{A_n\}_{n\in\mathbb{N}}$, the event $\cap_{n=1}^{\infty} \cup_{i=n}^{\infty} A_i$ is denoted by $\limsup_{n\to\infty} A_n$ (also called the event that A_n occurs infinitely often); the event $\cup_{n=1}^{\infty} \cap_{i=n}^{\infty} A_i$ is denoted by $\liminf_{n\to\infty} A_n$. A (real) random variable is a measurable function that maps Ω into $\mathbb{R}$. An extended (real) random variable can also take the values $\pm\infty$. The expected value of a random variable Y is $\mathbb{E}[Y] = \int_\Omega Y \, d\mathbb{P}$. A sequence of random variables $\{Y_n\}_{n\in\mathbb{N}}$ is said to converge surely to a random variable Y if $\lim_{n\to\infty} Y_n(\omega) = Y(\omega)$ for all $\omega \in \Omega$; the sequence is said to converge almost surely if $\mathbb{P}(\{\lim_{n\to\infty} Y_n = Y\}) = 1$. Finally, if $\varphi(\omega)$ is a property that is either true or false for a given $\omega \in \Omega$, the event that denotes the set of all samples ω for which $\varphi(\omega)$ holds, i.e., $\{\omega \in \Omega \mid \varphi(\omega) \text{ holds}\}$, is written as $\{\varphi\}$, e.g., $\{\omega \in \Omega \mid \lim_{n\to\infty} Y_n(\omega) = Y(\omega)\}$ is simply written as $\{\lim_{n\to\infty} Y_n = Y\}$. The Poisson random variable with parameter λ is denoted by $\mathrm{Poisson}(\lambda)$. The binomial random variable with parameters n and p is denoted by $\mathrm{Binomial}(n, p)$.

Let $f(n)$ and $g(n)$ be two functions with domain and range $\mathbb{N}$ or $\mathbb{R}$. The function $f(n)$ is said to be $O(g(n))$, denoted as $f(n) \in O(g(n))$, if there exists two constants M and n_0 such that $f(n) \leq M g(n)$ for all $n \geq n_0$. The function $f(n)$ is said to be $\Omega(g(n))$, denoted as $f(n) \in \Omega(g(n))$, if there exists constants M and n_0 such that $f(n) \geq M g(n)$ for all $n \geq n_0$. The function $f(n)$ is said to be $\Theta(g(n))$, denoted as $f(n) \in \Theta(g(n))$, if $f(n) \in O(g(n))$ and $f(n) \in \Omega(g(n))$.

Let $\mathcal{X}$ be a subset of $\mathbb{R}^d$. A (directed) graph $G = (V, E)$ on $\mathcal{X}$ is composed of a vertex set V and an edge set E, such that V is a finite subset of $\mathcal{X}$, and E is a subset of $V \times V$. A directed path on G is a sequence $(v_1, v_2, \ldots, v_n)$ of vertices such that $(v_i, v_{i+1}) \in E$ for all $1 \leq i \leq n - 1$. Given a vertex $v \in V$, the sets $\{u \in V \mid (u, v) \in E\}$ and $\{u \in V \mid (v, u) \in E\}$ are said to be its incoming neighbors and outgoing neighbors, respectively. A (directed) tree is a directed graph, in which each vertex but one has a unique incoming neighbor; the vertex with no incoming neighbor is called the root vertex. Vertices of a tree are often also called nodes.

B Proof of Theorem 33 (Non-optimality of RRT)

For simplicity, the theorem will be proven assuming that (i) the environment contains no obstacles, i.e., $\mathcal{X}_{\text{free}} = [0,1]^d$, and (ii) the parameter η of the steering procedure is set large enough, e.g., $\eta \geq \text{diam}\,(\mathcal{X}_{\text{free}}) = \sqrt{d}$. One one hand, considering this case is enough to prove that the RRT algorithm is not asymptotically optimal, as it demonstrates a case for which the RRT algorithm fails to converge to an optimal solution, although the problem instance is clearly robustly optimal. On the other hand, these assumptions are not essential, and the claims extend to the more general case, but the technical details of the proof are considerably more complicated.

The proof can be outlined as follows. Order the vertices in the RRT according to the iteration at which they are added to the tree. The set of vertices that contains the k-th child of the root along with all its descendants in the tree is called the k-th branch of the tree. First, it is shown that a necessary condition for the asymptotic optimality of RRT is that infinitely many branches of the tree contain vertices outside a small ball centered at the initial condition. Then, the RRT algorithm is shown to violate this condition, with probability one.

B.1 A necessary condition

First, we provide a necessary condition for the RRT algorithm to be asymptotically optimal.

Lemma 44 *Let $0 < R < \inf_{y \in \mathcal{X}_{\text{goal}}} \|y - x_{\text{init}}\|$. The event $\{\lim_{N\to\infty} Y_n^{\text{RRT}} = c^*\}$ occurs only if the k-th branch of the RRT contains vertices outside the R-ball centered at x_{init} for infinitely many k.*

Proof Let $\{x_1, x_2, \dots\}$ denote the set of children to the root vertex in the order they are added to the tree. Let $\Gamma(x_k)$ denote the optimal cost of a path starting from the root vertex, passing through x_k, and reaching the goal region. By our assumption that the measure of the set of all points that are on the optimal path is zero (see Assumption 27 and Lemma 28), the probability that $\Gamma(x_k) = c^*$ is zero for all $k \in \mathbb{N}$. Hence,

$$\mathbb{P}\Big(\bigcup\nolimits_{k\in\mathbb{N}} \{\Gamma(x_k) = c^*\}\Big) \;\leq\; \sum_{k=1}^{\infty} \mathbb{P}\big(\{\Gamma(x_k) = c^*\}\big) = 0.$$

Let A_k denote the event that at least one vertex in the k-th branch of the tree is outside the ball of radius R centered at x_{init} in some iteration of the RRT algorithm. Consider the case when the event $\{\limsup_{k\to\infty} A_k\}$ does not occur and the events $\{\Gamma(x_k) > c^*\}$ occur for all $k \in \mathbb{N}$. Then, A_k occurs for only finitely many k. Let K denote the largest number such that A_K occurs. Then, the cost of the best path in the tree is at least $\sup\{\Gamma(x_k)\,|\, k \in \{1, 2, \dots, K\}\}$, which is strictly larger than c^*, since $\{\Gamma(x_k) > c^*\}$ for all finite k. Thus, $\lim_{n\to\infty} Y_n^{\text{RRT}} > c^*$ must hold. That is, we have argued that

$$\Big(\limsup_{k\to\infty} A_k\Big)^c \cap \Big(\bigcap_{k\in\mathbb{N}} \{\Gamma(x_k) > c^*\}\Big) \subseteq \Big\{\lim_{n\to\infty} Y_n^{\text{RRT}} > c^*\Big\}.$$

Taking the complement of both sides and using monotonicity of probability measures,

$$\begin{aligned}\mathbb{P}\left(\{\lim_{n\to\infty} Y_n^{\text{RRT}} = c^*\}\right) \;&\leq\; \mathbb{P}\Big(\big(\limsup_{k\to\infty} A_k\big) \cup \big(\bigcup\nolimits_{k\in\mathbb{N}} \{\Gamma(x_k) = c^*\}\big)\Big),\\ &\leq\; \mathbb{P}\Big(\limsup_{k\to\infty} A_k\Big) + \mathbb{P}\Big(\bigcup\nolimits_{k\in\mathbb{N}} \{\Gamma(x_k) = c^*\}\Big),\end{aligned}$$

where the last inequality follows from the union bound. The lemma follows from the fact that the last term in the right hand side is equal to zero as shown above. □

B.2 Length of the first path in a branch

The following result provides a useful characterization of the RRT structure.

Lemma 45 *Let $U = \{X_1, X_2, \ldots, X_n\}$ be a set of independently sampled and uniformly distributed points in the d-dimensional unit cube, $[0,1]^d$. Let X_{n+1} be a point that is sampled independently from all the other points according to the uniform distribution on $[0,1]^d$. Then, the probability that among all points in U the point X_i is the one that is closest to X_{n+1} is $1/n$, for all $i \in \{1, 2, \ldots, n\}$. Moreover, the expected distance from X_{n+1} to its nearest neighbor in U is $n^{-1/d}$.*

Proof Since the probability distribution is uniform, the probability that X_{n+1} is closest to X_i is the same for all $i \in \{1, 2, \ldots, n\}$, which implies that this probability is equal to $1/n$. The expected distance to the closest point in U is an application of the order statistics of the uniform distribution. □

An immediate consequence of this result is that each vertex of the RRT has unbounded degree, almost surely, as the number of samples approaches infinity.

One can also define a notion of infinite paths in the RRT, as follows. Let Λ be the set of infinite sequences of natural numbers $\alpha = (\alpha_1, \alpha_2, \ldots)$. For any $i \in \mathbb{N}$, let $\pi_i : \Sigma \to \mathbb{N}^i, (\alpha_1, \alpha_2, \ldots, \alpha_i, \ldots) \mapsto (\alpha_1, \alpha_2, \ldots, \alpha_i)$, be a function returning the prefix of length i of an infinite sequence in Λ. The lexicographic ordering of Λ is such that, given $\alpha, \beta \in \Sigma$, $\alpha \leq \beta$ if and only if there exists $j \in \mathbb{N}$ such that $\alpha_i = \beta_i$ for all $i \in \mathbb{N}$, $i \leq j-1$, and $\alpha_j \leq \beta_j$. This is a total ordering of Λ, since $\mathbb{N}$ is a totally ordered set. Given $\alpha \in \Lambda$ and $i \in \mathbb{N}$, let $\mathcal{L}_{\pi_i(\alpha)}$ be the sum of the distances from the root vertex x_{init} to its α_1-th child, from this vertex to its α_2-th child, etc., for a total of i terms. Because of Lemma 45, this construction is well defined, almost surely, for a sufficiently large number of samples. For any infinite sequence $\alpha \in \Lambda$, let $\mathcal{L}_\alpha = \lim_{i \to +\infty} \mathcal{L}_{\pi_i(\alpha)}$; the limit exists since $\mathcal{L}_{\pi_i(\alpha)}$ is non-decreasing in i.

Consider infinite strings of the form $\mathbf{k} = (k, 1, 1, \ldots)$, $k \in \mathbb{N}$, and introduce the shorthand $\mathcal{L}_{\mathbf{k}} := \mathcal{L}_{(k,1,1,\ldots)}$. The following lemma shows that, for any $k \in \mathbb{N}$, $\mathcal{L}_{\mathbf{k}}$ has finite expectation, which immediately implies that $\mathcal{L}_{\mathbf{k}}$ takes only finite values with probability one. The lemma also provides a couple of other useful properties of $\mathcal{L}_{\mathbf{k}}$, which will be used later on.

Lemma 46 *The expected value $\mathbb{E}[\mathcal{L}_{\mathbf{k}}]$ is non-negative and finite, and monotonically non-increasing, in the sense that $\mathbb{E}[\mathcal{L}_{\mathbf{k+1}}] \leq \mathbb{E}[\mathcal{L}_{\mathbf{k}}]$, for any $k \in \mathbb{N}$. Moreover, $\lim_{k \to \infty} \mathbb{E}[\mathcal{L}_{\mathbf{k}}] = 0$.*

Proof Under the simplifying assumptions that there are no obstacles in the unit cube and η is large enough, the vertex set V_n^{RRT} of the graph maintained by the RRT algorithm is precisely the first n samples and each new sample is connected to its nearest neighbor in V_n^{RRT}.

Define Z_i as a random variable describing the contribution to $\mathcal{L}_{\mathbf{1}}$ realized at iteration i; in other words, Z_i is the distance of the i-th sample to its nearest neighbor among the first $i-1$ samples if the i-th sample is on the path used in computing $\mathcal{L}_{\mathbf{1}}$, and zero otherwise. Then, using Lemma 45,

$$\mathbb{E}[\mathcal{L}_{\mathbf{1}}] = \mathbb{E}\left[\sum_{i=1}^{\infty} Z_i\right] = \sum_{i=1}^{\infty} \mathbb{E}[Z_i] = \sum_{i=1}^{\infty} i^{-1/d}\, i^{-1} = \mathtt{Zeta}(1 + 1/d),$$

where the second equality follows from the monotone convergence theorem and `Zeta` is the Riemann zeta function. Since $\mathtt{Zeta}(y)$ is finite for any $y > 1$, $\mathbb{E}[\mathcal{L}_{\mathbf{1}}]$ is a finite number for all $d \in \mathbb{N}$.

Let N_k be the iteration at which the first sample contributing to $\mathcal{L}_k$ is generated. Then, an argument similar to the one given above yields

$$\mathbb{E}[\mathcal{L}_{\mathbf{k+1}}] = \sum_{i=N_k+1}^{\infty} i^{-(1+1/d)} = \mathbb{E}[\mathcal{L}_{\mathbf{1}}] - \sum_{i=1}^{N_k} i^{-(1+1/d)}.$$

Then, clearly, $\mathbb{E}[\mathcal{L}_{\mathbf{k+1}}] < \mathbb{E}[\mathcal{L}_{\mathbf{k}}]$ for all $k \in \mathbb{N}$. Moreover, since $N_k \geq k$, it is the case that $\lim_{k\to\infty} \mathbb{E}[\mathcal{L}_{\mathbf{k}}] = 0$. □

B.3 Length of the longest path in a branch

Given $k \in \mathbb{N}$, and the sequence $\mathbf{k} = (k, 1, 1, \ldots)$, the quantity $\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha$ is an upper bound on the length of any path in the k-th branch of the RRT, or in any of the following branches. The next result bounds the probability that this quantity is very large.

Lemma 47 *For any* $\epsilon > 0$,

$$\mathbb{P}\left(\left\{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > \epsilon\right\}\right) \leq \frac{\mathbb{E}[\mathcal{L}_{\mathbf{k}}]}{\epsilon}.$$

First, we state and prove the following intermediate result.

Lemma 48 $\mathbb{E}[\mathcal{L}_\alpha] \leq \mathbb{E}[\mathcal{L}_{\mathbf{k}}]$, *for all* $\alpha \geq \mathbf{k}$.

Proof The proof is by induction. Since $\alpha \geq \mathbf{k}$, then $\pi_1(\alpha) \geq k$, and Lemma 46 implies that $\mathbb{E}[\mathcal{L}_{(\pi_1(\alpha),1,1,\ldots)}] \leq \mathbb{E}[\mathcal{L}_{\mathbf{k}}]$. Moreover, it is also the case that, for any $i \in \mathbb{N}$ (and some abuse of notation), $\mathbb{E}[\mathcal{L}_{(\pi_{i+1}(\alpha),1,1,\ldots)}] \leq \mathbb{E}[\mathcal{L}_{(\pi_i(\alpha),1,1,\ldots)}]$, by a similar argument considering a tree rooted at the last vertex reached by the finite path $\pi_i(\alpha)$. Since $(\pi_{i+1}(\alpha), 1, 1, \ldots) \geq (\pi_i(\alpha), 1, 1, \ldots) \geq (k, 1, 1, \ldots)$, the result follows. □

Proof of Lemma 47 Define the random variable $\bar{\alpha} := \inf\{\alpha \geq \mathbf{k} \,|\, \mathcal{L}_\alpha > \epsilon\}$, and set $\bar{\alpha} := \mathbf{k}$ if $\mathcal{L}_\alpha \leq \epsilon$ for all $\alpha \geq \mathbf{k}$. Note that $\bar{\alpha} \geq \mathbf{k}$ holds surely. Hence, by Lemma 48, $\mathbb{E}[\mathcal{L}_{\bar{\alpha}}] \leq \mathbb{E}[\mathcal{L}_{\mathbf{k}}]$. Let I_ϵ be the indicator random variable for the event $S_\epsilon := \{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > \epsilon\}$. Then,

$$\mathbb{E}[\mathcal{L}_{\mathbf{k}}] \geq \mathbb{E}[\mathcal{L}_{\bar{\alpha}}] = \mathbb{E}[\mathcal{L}_{\bar{\alpha}} I_\epsilon] + \mathbb{E}[\mathcal{L}_{\bar{\alpha}}(1 - I_\epsilon)] \geq \epsilon\, \mathbb{P}(S_\epsilon),$$

where the last inequality follows from the fact that $\mathcal{L}_{\bar{\alpha}}$ is at least ϵ whenever the event S_ϵ occurs. □

A useful corollary of Lemmas 46 and 47 is the following.

Corollary 49 *For any* $\epsilon > 0$, $\lim_{k\to\infty} \mathbb{P}(\{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > \epsilon\}) = 0$.

B.4 Violation of the necessary condition

Recall from Lemma 44 that a necessary condition for asymptotic optimality is that the k-th branch of the RRT contains vertices outside the R-ball centered at x_{init} for infinitely many k, where $0 < R < \inf_{y \in \mathcal{X}_{\text{goal}}} \|y - x_{\text{init}}\|$. Clearly, the latter event can occur only if longest path in the k-th branch of the RRT is longer than R for infinitely many k. That is,

$$\mathbb{P}\left(\left\{\lim_{n\to\infty} Y_n^{\text{RRT}} = c^*\right\}\right) \leq \mathbb{P}\left(\limsup_{k\to\infty} \left\{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > R\right\}\right).$$

The event on the right hand side is monotonic in the sense that $\{\sup_{\alpha > \mathbf{k}+1} \mathcal{L}_\alpha > R\} \supseteq \{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > R\}$ for all $k \in \mathbb{N}$. Hence, $\lim_{k\to\infty}\{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > R\}$ exists. In particular, $\mathbb{P}(\limsup_{k\to\infty}\{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > R\}) = \mathbb{P}(\lim_{k\to\infty}\{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > R\}) = \lim_{k\to\infty} \mathbb{P}(\{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > R\})$, where the last equality follows from the continuity of probability measures. Since $\lim_{k\to\infty} \mathbb{P}\left(\{\sup_{\alpha \geq \mathbf{k}} \mathcal{L}_\alpha > R\}\right) = 0$ for all $R > 0$ by Corollary 49, $\mathbb{P}(\{\lim_{n\to\infty} Y_n^{\mathrm{RRT}} = c^*\}) = 0$.

C Proof of Theorem 34 (Asymptotic optimality of PRM*)

An outline of the proof is given below, before the details are provided.

C.1 Outline of the proof

Let σ^* denote a robustly optimal path. By definition, σ^* has weak δ-clearance. First, define a sequence $\{\delta_n\}_{n\in\mathbb{N}}$ such that $\delta_n > 0$ for all $n \in \mathbb{N}$ and δ_n approaches zero as n approaches infinity. Construct a sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths such that σ_n has strong δ_n-clearance for all $n \in \mathbb{N}$ and σ_n converges to σ^* as n approaches infinity.

Second, define a sequence $\{q_n\}_{n\in\mathbb{N}}$. For all $n \in \mathbb{N}$, construct a set $B_n = \{B_{n,1}, B_{n,2}, \ldots, B_{n,M_n}\}$ of overlapping balls, each with radius q_n, that collectively "cover" the path σ_n. See Figures 23 and 24. Let $x_m \in B_{n,m}$ and $x_{m+1} \in B_{n,m+1}$ be any two points from two consecutive balls in B_n. Construct B_n such that (i) x_m and x_{m+1} have distance no more than the connection radius $r(n)$ and (ii) the straight path connecting x_m and x_{m+1} lies entirely within the obstacle free space. These requirements can be satisfied by setting δ_n and q_n to certain constant fractions of $r(n)$.

Let A_n denote the event that each ball in B_n contains at least one vertex of the graph returned by the PRM* algorithm, when the algorithm is run with n samples. Third, show that A_n occurs for all large n, with probability one. Clearly, in this case, the PRM* algorithm will connect the vertices in consecutive balls with an edge, and any path formed in this way will be collision-free.

Finally, show that any sequence of paths generated in this way converges to the optimal path σ^*. Using the robustness of σ^*, show that the cost of the best path in the graph returned by the PRM* algorithm converges to $c(\sigma^*)$ almost surely.

C.2 Construction of the sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths

The following lemma establishes a connection between the notions of strong and weak δ-clearance.

Lemma 50 *Let σ^* be a path be a path that has strong δ-clearance. Let $\{\delta_n\}_{n\in\mathbb{N}}$ be a sequence of real numbers such that $\lim_{n\to\infty} \delta_n = 0$ and $0 \leq \delta_n \leq \delta$ for all $n \in \mathbb{N}$. Then, there exists a sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths such that $\lim_{n\to\infty} \sigma_n = \sigma^*$ and σ_n has strong δ_n-clearance for all $n \in \mathbb{N}$.*

Proof First, define a sequence $\{\mathcal{X}_n\}_{n\in\mathbb{N}}$ of subsets of $\mathcal{X}_{\mathrm{free}}$ such that $\mathcal{X}_n$ is the closure of the δ_n-interior of $\mathcal{X}_{\mathrm{free}}$, i.e.,

$$\mathcal{X}_n := \mathrm{cl}(\mathrm{int}_{\delta_n}(\mathcal{X}_{\mathrm{free}}))$$

for all $n \in \mathbb{N}$. Note that, by definition, (i) $\mathcal{X}_n$ are closed subsets of $\mathcal{X}_{\mathrm{free}}$, and (ii) any point $\mathcal{X}_n$ has distance at least δ_n to any point in the obstacle set $\mathcal{X}_{\mathrm{obs}}$.

Then, construct the sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths, where $\sigma_n \in \Sigma_{\mathcal{X}_n}$, as follows. Let $\psi : [0,1] \to \Sigma_{\mathrm{free}}$ denote the homotopy with $\psi(0) = \sigma^*$; the existence of ψ is guaranteed by weak δ-clearance of σ^*. Define

$$\alpha_n := \max_{\alpha\in[0,1]} \{\alpha \,|\, \psi(\alpha) \in \Sigma_{\mathcal{X}_n}\} \quad \text{and} \quad \sigma_n := \psi(\alpha_n).$$

Since $\Sigma_{\mathcal{X}_n}$ is closed, the maximum in the definition of α_n is attained. Moreover, since $\psi(1)$ has strong δ-clearance and $\delta_n \leq \delta$, $\sigma_n \in \Sigma_{\mathcal{X}_n}$, which implies the strong δ_n-clearance of σ_n.

Clearly, $\bigcup_{n\in\mathbb{N}} \mathcal{X}_n = \mathcal{X}_{\text{free}}$, since $\lim_{n\to\infty} \delta_n = 0$. Also, by weak δ-clearance of σ^*, for any $\alpha \in (0,1]$, there exists some $\delta_\alpha \in (0,\delta]$ such that $\psi(\alpha)$ has strong δ_α-clearance. Then, $\lim_{n\to\infty} \alpha_n = 0$, which implies $\lim_{n\to\infty} \sigma_n = \sigma^*$. □

Recall that the connection radius of the PRM* algorithm was defined as

$$r_n = \gamma_{\text{PRM}} \left(\frac{\log n}{n}\right)^{1/d} > 2(1+1/d)^{1/d} \left(\frac{\mu(X_{\text{free}})}{\zeta_d}\right)^{1/d} \left(\frac{\log n}{n}\right)^{1/d}$$

(see Algorithm 4 and the definition of the `Near` procedure in Section 3.1). Let θ_1 be a small positive constant; the precise value of θ_1 will be provided shortly in the proof of Lemma 52. Define

$$\delta_n := \min\left\{\delta, \frac{1+\theta_1}{2+\theta_1} r_n\right\}, \qquad \text{for all } n \in \mathbb{N}.$$

By definition, $0 \leq \delta_n \leq \delta$ holds. Moreover, $\lim_{n\to\infty} \delta_n = 0$, since $\lim_{n\to\infty} r_n = 0$. Then, by Lemma 50, there exists a sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths such that $\lim_{n\to\infty} \sigma_n = \sigma^*$ and σ_n has strong δ_n-clearance for all $n \in \mathbb{N}$.

C.3 Construction of the sequence $\{B_n\}_{n\in\mathbb{N}}$ of sets of balls

First, a construction of a finite set of balls that collectively "cover" a path σ_n is provided. The construction is illustrated in Figure 23.

Definition 51 (Covering balls) *Given a path $\sigma_n : [0,1] \to \mathcal{X}$, and the real numbers $q_n, l_n \in \mathbb{R}_{>0}$, the set* `CoveringBalls`$(\sigma_n, q_n, l_n)$ *is defined as a set $\{B_{n,1}, B_{n,2}, \ldots, B_{n,M_n}\}$ of M_n balls of radius q_n such that $B_{n,m}$ is centered at $\sigma(\tau_m)$, and*

- *the center of $B_{n,1}$ is $\sigma(0)$, i.e., $\tau_1 = 0$,*
- *the centers of two consecutive balls are exactly l_n apart, i.e., $\tau_m := \min\{\tau \in [\tau_{m-1}, 1] \mid \|\sigma(\tau) - \sigma(\tau_{m-1})\| \geq l_n\}$ for all $m \in \{2, 3, \ldots, M_n\}$,*
- *and $M-1$ is the largest number of balls that can be generated in this manner while the center of the last ball, B_{n,M_n} is $\sigma(1)$, i.e., $\tau_{M_n} = 1$.*

For each $n \in \mathbb{N}$, define

$$q_n := \frac{\delta_n}{1+\theta_1}.$$

Construct the set $B_n = \{B_{n,1}, B_{n,2}, \ldots, B_{n,M_n}\}$ of balls as $B_n :=$ `CoveringBalls`$(\sigma_n, q_n, \theta_1 q_n)$ using Definition 51 (see Figure 23). By construction, each ball in B_n has radius q_n and the centers of consecutive balls in B_n are $\theta_1 q_n$ apart (see Figure 24 for an illustration of covering balls with this set of parameters). The balls in B_n collectively cover the path σ_n.

C.4 The probability that each ball in B_n contains at least one vertex

Recall that $G_n^{\text{PRM}^*} = (V_n^{\text{PRM}^*}, E_n^{\text{PRM}^*})$ denotes the graph returned by the PRM* algorithm, when the algorithm is run with n samples. Let $A_{n,m}$ denote the event that the ball $B_{n,m}$ contains at least one vertex of the graph generated by the PRM* algorithm, i.e., $A_{n,m} = \{B_{n,m} \cap V_n^{\text{PRM}^*} \neq \emptyset\}$. Let A_n denote the event that all balls in B_n contain at least one vertex of the PRM* graph, i.e., $A_n = \bigcap_{m=1}^{M_n} A_{n,m}$.

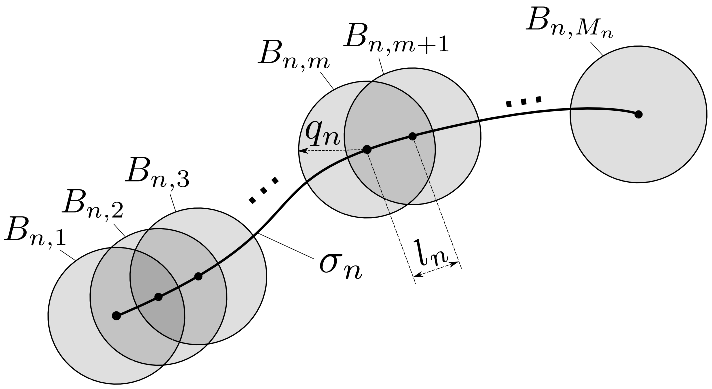

Figure 23: An illustration of the `CoveringBalls` construction. A set of balls that collectively cover the trajectory σ_n is shown. All balls have the same radius, q_n. The spacing between the centers of two consecutive balls is l_n.

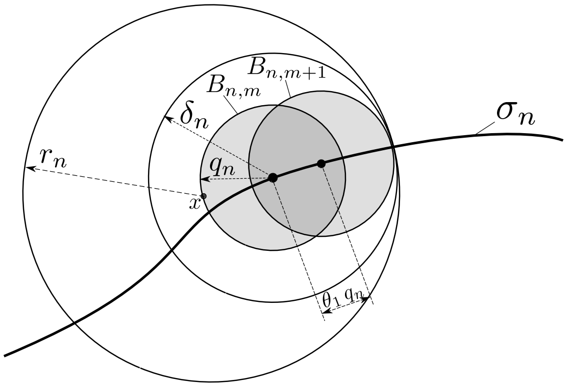

Figure 24: An illustration of the covering balls for PRM* algorithm. The δ_n-ball is guaranteed to be inside the obstacle-free space. The connection radius r_n is also shown as the radius of the connection ball centered at a vertex $x \in B_{n,m}$. The vertex x is connected to all other vertices that lie within the connection ball.

Lemma 52 *If* $\gamma_{\mathrm{PRM}} > 2\,(1+1/d)^{1/d} \left(\frac{\mu(X_{\mathrm{free}})}{\zeta_d}\right)^{1/d}$*, then there exists a constant* $\theta_1 > 0$ *such that the event that every ball in* B_n *contains at least one vertex of the PRM* graph occurs for all large enough* n *with probability one, i.e.,*

$$\mathbb{P}\left(\liminf_{n\to\infty} A_n\right) = 1.$$

Proof The proof is based on a Borel-Cantelli argument which can be summarized as follows. Recall that A_n^c denotes the complement of A_n. First, the sum $\sum_{n=1}^{\infty} \mathbb{P}(A_n^c)$ is shown to be bounded. By the Borel-Cantelli lemma (Grimmett and Stirzaker, 2001), this implies that the probability that A_n holds infinitely often as n approaches infinity is zero. Hence, the probability that A_n holds infinitely often is one. In the rest of the proof, an upper bound on $\mathbb{P}(A_n)$ is computed, and this upper bound is shown to be summable.

First, compute a bound on the number of balls in B_n as follows. Let s_n denote the length of σ_n, i.e., $s_n := \mathrm{TV}(\sigma_n)$. Recall that the balls in B_n were constructed such that the centers of two consecutive balls in B_n have distance $\theta_1\, q_n$. The segment of σ_n that starts at the center of $B_{n,m}$ and ends at the center of $B_{n,m+1}$ has length at least $\theta_1 q_n$, except for the last segment, which has length less than or equal to $\theta_1 q_n$. Let $n_0 \in \mathbb{N}$ be the number such that $\delta_n < \delta$ for all $n \geq n_0$. Then, for all $n \geq n_0$,

$$\begin{aligned} \operatorname{card}(B_n) = M_n \quad &\leq \quad \frac{s_n}{\theta_1 q_n} = \frac{(1+\theta_1)s_n}{\theta_1 \delta_n} = \frac{(2+\theta_1)\, s_n}{\theta_1\, r_n} \\ &= \quad \frac{(2+\theta_1)\, s_n}{\theta_1\, \gamma_{\mathrm{PRM}}} \left(\frac{n}{\log n}\right)^{1/d}. \end{aligned}$$

Second, compute the volume of a single ball in B_i as follows. Recall that $\mu(\cdot)$ denotes the usual Lebesgue measure, and ζ_d denotes the volume of a unit ball in the d-dimensional Euclidean space. For all $n \geq n_0$,

$$\mu(B_{n,m}) = \zeta_d\, q_n^d = \zeta_d \left(\frac{\delta_n}{1+\theta_1}\right)^d = \zeta_d \left(\frac{r_n}{2+\theta_1}\right)^d = \zeta_d \left(\frac{\gamma_{\mathrm{PRM}}}{2+\theta_1}\right)^d \frac{\log n}{n}$$

For all $n \geq I$, the probability that a single ball, say $B_{n,1}$, does not contain a vertex of the graph generated by the PRM* algorithm, when the algorithm is run with n samples, is

$$\begin{aligned} \mathbb{P}\left(A_{n,1}^c\right) \quad &= \quad \left(1 - \frac{\mu(B_{n,1})}{\mu(X_{\mathrm{free}})}\right)^n \\ &= \quad \left(1 - \frac{\zeta_d}{\mu(X_{\mathrm{free}})} \left(\frac{\gamma_{\mathrm{PRM}}}{2+\theta_1}\right)^d \frac{\log n}{n}\right)^n \end{aligned}$$

Using the inequality $(1 - 1/f(n))^r \leq e^{-r/f(n)}$, the right-hand side can be bounded as

$$\mathbb{P}(A_{n,1}) \leq e^{-\frac{\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{PRM}}}{2+\theta_1}\right)^d \log n} = n^{-\frac{\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{PRM}}}{2+\theta_1}\right)^d}.$$

Hence,

$$\begin{aligned} \mathbb{P}\left(A_n^c\right) = \mathbb{P}\left(\bigcup_{m=1}^{M_n} A_{n,m}^c\right) \quad &\leq \quad \sum_{m=1}^{M_n} \mathbb{P}\left(A_{n,m}^c\right) = M_n\, \mathbb{P}(A_{n,1}^c) \\ &\leq \quad \frac{(2+\theta_1)s_n}{\theta_1\, \gamma_{\mathrm{PRM}}} \left(\frac{n}{\log n}\right)^{1/d} i^{-\frac{\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{PRM}}}{2+\theta_1}\right)^d} \\ &= \quad \frac{(2+\theta_1)s_n}{\theta_1\, \gamma_{\mathrm{PRM}}} \frac{1}{(\log n)^d}\; n^{-\left(\frac{\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{PRM}}}{2+\theta_1}\right)^d - \frac{1}{d}\right)} \end{aligned}$$

where the first inequality follows from the union bound.

Finally, $\sum_{n=1}^{\infty} \mathbb{P}(A_n^c) < \infty$ holds, if $\frac{\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{PRM}}}{2+\theta_1}\right)^d - \frac{1}{d} > 1$, which can be satisfied for any $\gamma_{PRM} > 2(1+1/d)^{1/d}\left(\frac{\mu(X_{\mathrm{free}})}{\zeta_d}\right)^{1/d}$ by appropriately choosing θ_1. Then, by the Borel-Cantelli lemma (Grimmett and Stirzaker, 2001), $\mathbb{P}(\limsup_{n\to\infty} A_n^c) = 0$, which implies $\mathbb{P}(\liminf_{n\to\infty} A_n) = 1$. $\square$

C.5 Connecting the vertices in subsequent balls in B_n

Let $Z_n := \{x_1, x_2, \ldots, x_{M_n}\}$ be any set of points such that $x_m \in B_{n,m}$ for each $m \in \{1, 2, \ldots, M_n\}$. The following lemma states that for all $n \in \mathbb{N}$ and all $m \in \{1, 2, \ldots, M_n - 1\}$, the distance between x_m and x_{m+1} is less than the connection radius, r_n, which implies that the PRM* algorithm will attempt to connect the two points x_m and x_{m+1} if they are in the vertex set of the PRM* algorithm.

Lemma 53 *If $x_{n,m} \in B_{n,m}$ and $x_{n,m+1} \in B_{n,m+1}$, then $\|x_{n,m+1} - x_{n,m}\| \leq r_n$, for all $n \in \mathbb{N}$ and all $m \in \{1, 2, \ldots, M_i - 1\}$.*

Proof Recall that each ball in B_n has radius $q_n = \frac{\delta_n}{(1+\theta_1)}$. Given any two points $x_m \in B_{n,m}$ and $x_{m+1} \in B_{n,m+1}$, all of the following hold: (i) x_m has distance q_n to the center of $B_{n,m}$, (ii) x_{m+1} has distance q_n to the center of $B_{n,m+1}$, and (iii) centers of $B_{n,m}$ and $B_{n,m+1}$ have distance $\theta_1 \, q_n$ to each other. Then,

$$\|x_{n,m+1} - x_{n,m}\| \leq (2 + \theta_1) \, q_n = \frac{2 + \theta_1}{1 + \theta_1} \, \delta_n \leq r_n,$$

where the first inequality is obtained by an application of the triangle inequality and the last inequality follows from the definition of $\delta_n = \min\{\delta, \frac{1+\theta_1}{2+\theta_1} \, r_n\}$. ☐

By Lemma 53, conclude that the PRM* algorithm will attempt to connect any two vertices in consecutive balls in B_n. The next lemma shows that any such connection attempt will, in fact, be successful. That is, the path connecting $x_{n,m}$ and $x_{n,m+1}$ is collision-free for all $m \in \{1, 2, \ldots, M_n\}$.

Lemma 54 *For all $n \in \mathbb{N}$ and all $m \in \{1, 2, \ldots, M_n\}$, if $x_m \in B_{n,m}$ and $x_{m+1} \in B_{n,m+1}$, then the line segment connecting $x_{n,m}$ and $x_{n,m+1}$ lies in the obstacle-free space, i.e.,*

$$\alpha \, x_{n,m} + (1 - \alpha) \, x_{n,m+1} \in X_{\text{free}}, \qquad \textit{for all } \alpha \in [0, 1].$$

Proof Recall that σ_n has strong δ_n-delta clearance and that the radius q_n of each ball in B_n was defined as $q_n = \frac{\delta_n}{1+\theta_1}$, where $\theta_1 > 0$ is a constant. Hence, any point along the trajectory σ_n has distance at least $(1 + \theta_1) \, q_n$ to any point in the obstacle set. Let y_m and y_{m+1} denote the centers of the balls $B_{n,m}$ and $B_{n,m+1}$, respectively. Since $y_m = \sigma(\tau_m)$ and $y_{m+1} = \sigma(\tau_{m+1})$ for some τ_m and τ_{m+1}, y_m and y_{m+1} also have distance $(1 + \theta_1) q_n$ to any point in the obstacle set.

Clearly, $\|x_m - y_m\| \leq q_n$. Moreover, the following inequality holds:

$$\|x_{m+1} - y_m\| \leq \|(x_m - y_{m+1}) + (y_{m+1} - y_m)\| \leq \|x_{m+1} - y_{m+1}\| + \|y_{m+1} - y_m\| \leq q_n + \theta_1 \, q_n = (1 + \theta_1) \, q_n.$$

where the second inequality follows from the triangle inequality and the third inequality follows from the construction of balls in B_n.

For any convex combination $x_\alpha := \alpha \, x_m + (1 - \alpha) \, x_{m+1}$, where $\alpha \in [0, 1]$, the distance between x_α and y_m can be bounded as follows:

$$\begin{aligned} \big\| \big(\alpha \, x_m + (1 + \alpha) \, x_{m+1}\big) - y_m \big\| &= \big\| \alpha \, (x_m - y_m) + (1 + \alpha) \, (x_{m+1} - y_m) \big\| \\ &= \alpha \, \|x_m - y_m\| + (1 + \alpha) \, \|x_{m+1} - y_m\| \\ &= \alpha \, q_n + (1 + \alpha) \, (1 + q_n) \leq (1 + \theta_1) \, q_n, \end{aligned}$$

where the second equality follows from the linearity of the norm. Hence, any point along the line segment connecting x_m and x_{m+1} has distance at most $(1 + \theta_1) \, q_n$ to y_m. Since, y_m has distance at least $(1 + \theta_1) q_n$ to any point in the obstacle set, the line segment connecting x_m and x_{m+1} is collision-free. ☐

C.6 Convergence to the optimal path

Let P_n denote the set of all paths in the graph $G_n^{\mathrm{PRM}^*} = (V_n^{\mathrm{PRM}^*}, E_n^{\mathrm{PRM}^*})$. Let σ'_n be the path that is closest to σ_n in terms of the bounded variation norm among all those paths in P_n, i.e., $\sigma'_n := \min_{\sigma' \in P_n} \|\sigma' - \sigma_n\|$. Note that the sequence $\{\sigma'_n\}_{n\in\mathbb{N}}$ is a random sequence of paths, since the graph $G_n^{\mathrm{PRM}^*}$, hence the set P_n of paths is random. The following lemma states that the bounded variation distance between σ'_n and σ_n approaches to zero, with probability one.

Lemma 55 *The random variable* $\|\sigma'_n - \sigma_n\|_{\mathrm{BV}}$ *converges to zero almost surely, i.e.,*

$$\mathbb{P}\left(\left\{\lim_{n\to\infty} \|\sigma'_n - \sigma_n\|_{\mathrm{BV}} = 0\right\}\right) = 1.$$

Proof The proof of this lemma is based on a Borel-Cantelli argument. It is shown that $\sum_{n\in\mathbb{N}} \mathbb{P}(\|\sigma'_n - \sigma_n\|_{\mathrm{BV}} > \epsilon)$ is finite for any $\epsilon > 0$, which implies that $\|\sigma'_n - \sigma_n\|$ converges to zero almost surely by the Borel-Cantelli lemma (Grimmett and Stirzaker, 2001). This proof uses a Poissonization argument in one of the intermediate steps. That is, a particular result is shown to hold in the Poisson process described in Lemma 11. Subsequently, the result is de-Poissonized, i.e., shown to hold also for the original process.

Fix some $\epsilon > 0$. Let $\alpha, \beta \in (0,1)$ be two constants, both independent of n. Recall that q_n is the radius of each ball in the set B_n of balls covering the path σ_n. Let $I_{n,m}$ denote the indicator variable for the event that the ball $B_{n,m}$ has no point that is within a distance $\beta\, q_n$ from the center of $B_{n,m}$. For a more precise definition, let $\beta\, B_{n,m}$ denote the ball that is centered at the center of $B_{n,m}$ and has radius $\beta\, r_n$. Then,

$$I_{n,m} := \begin{cases} 1, & \text{if } (\beta\, B_{n,m}) \cap V^{\mathrm{PRM}^*} = \emptyset, \\ 0, & \text{otherwise.} \end{cases}$$

Let K_n denote the number of balls in B_n that do not contain a vertex that is within a $\beta\, q_n$ distance to the center of that particular ball, i.e., $K_n := \sum_{m=1}^{M_n} I_{n,m}$.

Consider the event that $I_{n,m}$ holds for at most an α fraction of the balls in B_n, i.e., $\{K_n \leq \alpha\, M_n\}$. This event is important for the following reason. Recall that the vertices in subsequent balls in B_n are connected by edges in $G_n^{\mathrm{PRM}^*}$ by Lemmas 53 and 54. If only at most an α fraction of the balls do not have a vertex that is less than a distance of $\beta\, r_n$ from their centers (hence, a $(1-\alpha)$ fraction have at least one vertex within a distance of $\beta\, r_n$ from their centers), i.e., $\{K_n \leq \alpha\, M_n\}$ holds, then the bounded variation difference between σ'_n and σ_n is at most $(\sqrt{2}\,\alpha + \beta(1-\alpha))L \leq \sqrt{2}\,(\alpha+\beta)L$, where L is a finite bound on the length of all paths in $\{\sigma_n\}_{n\in\mathbb{N}}$, i.e., $L := \sup_{n\in\mathbb{N}} \mathrm{TV}(\sigma_n)$. That is,

$$\{K_n \leq \alpha\, M_n\} \subseteq \left\{\|\sigma'_n - \sigma_n\|_{\mathrm{BV}} \leq \sqrt{2}\,(\alpha+\beta)\,L\right\}$$

Taking the complement of both sides and using the monotonicity of probability measures,

$$\mathbb{P}\left(\left\{\|\sigma'_n - \sigma_n\|_{\mathrm{BV}} > \sqrt{2}\,(\alpha+\beta)\,L\right\}\right) \leq \mathbb{P}\left(\{K_n \geq \alpha\, M_n\}\right).$$

In the rest of the proof, it is shown that the right hand side of the inequality above is summable for all small $\alpha, \beta > 0$, which implies that $\mathbb{P}\left(\{\|\sigma'_n - \sigma_n\,\| > \epsilon\}\right)$ is summable for all small $\epsilon > 0$.

For this purpose, the process that provides independent uniform samples from $\mathcal{X}_{\mathrm{free}}$ is approximated by an equivalent Poisson process described in Section 2.2. A more precise definition is given as follows. Let $\{X_1, X_2, \ldots, X_n\}$ denote the binomial point process corresponding to the `SampleFree` procedure. Let $\nu < 1$ be a constant independent of n. Recall that $\mathrm{Poisson}(\nu\, n)$

denotes the Poisson random variable with intensity $\nu\, n$ (hence, mean value $\nu\, n$). Then, the process $\mathcal{P}_{\nu\, n} := \{X_1, X_2, \ldots, X_{\mathrm{Poisson}(\nu\, n)}\}$ is a Poisson process restricted to $\mu(\mathcal{X}_{\mathrm{free}})$ with intensity $\nu\, n/\mu(\mathcal{X}_{\mathrm{free}})$ (see Lemma 11). Thus, the expected number of points of this Poisson process is $\nu\, n$.

Clearly, the set of points generated by one process is a subset of the those generated by the other. However, since $\nu < 1$, in most trials the Poisson point process $\mathcal{P}_{\nu\, n}$ is a subset of the binomial point process.

Define the random variable $\widetilde{K}_n$ denote the number of balls of that fail to have one sample within a distance $\beta\, r_n$ to their centers, when the underlying point process is $\mathcal{P}_{\nu\, n}$ (instead of the independent uniform samples provided by the `SampleFree` procedure). In other words, $\widetilde{K}_n$ is the random variable that is defined similar to K_n, except that the former is defined with respect to the points of $\mathcal{P}_{\nu\, n}$ whereas the latter is defined with respect to the n samples returned by `SampleFree` procedure.

Since $\{\widetilde{K}_n > \alpha\, M_n\}$ is a decreasing event, i.e., the probability that it occurs increases if $\mathcal{P}_{\nu n}$ includes fewer samples, the following bound holds (see, e.g., Penrose, 2003)

$$\mathbb{P}\big(\, \{K_n \geq \alpha\, M_n\}\,\big) \leq \mathbb{P}\big(\{\widetilde{K}_n \geq \alpha\, M_n\}\big) + \mathbb{P}(\{\mathrm{Poisson}(\nu\, n) \geq n\}).$$

Since a Poisson random variable has exponentially-decaying tails, the second term on the right hand side can be bounded as

$$\mathbb{P}(\{\mathrm{Poisson}(\nu\, n) \geq n\}) \leq e^{-cn},$$

where $c > 0$ is a constant.

The first term on the right hand side can be computed directly as follows. First, for all small β, the balls of radius $\beta\, r_n$ are all disjoint (see Figure 25). Denote this set of balls by $\widetilde{B}_{n,m} = \{\widetilde{B}_{n,1}, \widetilde{B}_{n,2}, \ldots, \widetilde{B}_{n,M_n}\}$. More precisely, $\widetilde{B}_{n,m}$ is the ball of radius $\beta\, q_n$ centered at the center of $B_{n,m}$. Second, observe that the event $\{K_n > \alpha\, M_n\}$ is equivalent to the event that at least an α fraction of all the balls in $\widetilde{B}_n$ include at least one point of the process $\mathcal{P}_{\nu\, n}$. Since, the point process $\mathcal{P}_{\nu\, n}$ is Poisson and the balls in $\widetilde{B}_n$ are disjoint for all small enough β, the probability that a single ball in $\widetilde{B}_n$ does not contain a sample is $p_n := \exp(-\zeta_d\, (\beta q_n)^d\, \nu\, n/\mu(\mathcal{X}_{\mathrm{free}})) \leq \exp(-c\, \beta\, \nu\, \log n)$ for some constant c. Third, by the independence property of the Poisson point process, the number of balls in $\widetilde{B}_n$ that do not include a point of the point process $\mathcal{P}_{\nu\, n}$ is a binomial random variable with parameters M_n and p_n. Then, for all large n,

$$\mathbb{P}\Big(\big\{\widetilde{K}_n \geq \alpha\, M_n\big\}\Big) \leq \mathbb{P}\left(\{\mathrm{Binomial}(M_n, p_n) \geq \alpha M_n\}\right) \leq \exp(-M_n\, p_n).$$

Combining the two inequalities above, the following bound is obtained for the original sampling process

$$\mathbb{P}\left(\{K_n \geq \alpha\, M_n\}\right) \leq e^{-c\, n} + e^{-M_n\, p_n}.$$

Summing up both sides,

$$\sum_{n=1}^{\infty} \mathbb{P}\left(\{K_n \geq \alpha\, n\}\right) < \infty.$$

This argument holds for all $\alpha, \beta, \nu > 0$. Hence, for all $\epsilon > 0$,

$$\sum_{n=1}^{\infty} \mathbb{P}\left(\left\{\|\sigma'_n - \sigma_n\|_{\mathrm{BV}} > \epsilon\right\}\right) < \infty.$$

Then, by the Borel-Cantelli lemma, $\mathbb{P}\left(\{\lim_{n\to\infty} \|\sigma'_n - \sigma_n\|_{\mathrm{BV}} = 0\}\right) = 1$. □

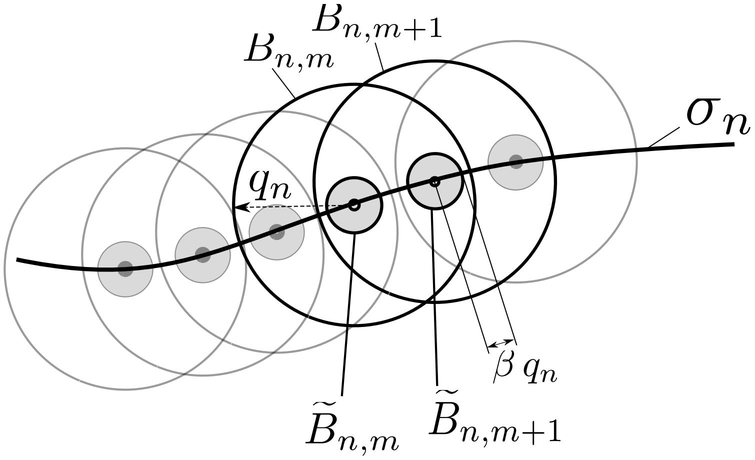

Figure 25: The set $\widetilde{B}_{n,m}$ of non-intersection balls is illustrated.

Finally, the following lemma states that the cost of the minimum cost path in the graph returned by the PRM* algorithm converges to the optimal cost c^* with probability one. Recall that $Y_n^{\mathrm{PRM}^*}$ denotes the cost of the minimum-cost path in the graph returned by the PRM* algorithm, when the algorithm is run with n samples.

Lemma 56 *Under the assumptions of Theorem 34, the cost of the minimum-cost path present in the graph returned by the* PRM* *algorithm converges to the optimal cost c^* as the number of samples approaches infinity, with probability one, i.e.,*

$$\mathbb{P}\left(\left\{\lim_{n\to\infty} Y_n^{\mathrm{PRM}^*} = c^*\right\}\right) = 1.$$

Proof Recall that σ^* denotes the optimal path, and that $\lim_{n\to\infty}\sigma_n = \sigma^*$ holds surely. By Lemma 55, $\lim_{n\to\infty}\|\sigma'_n - \sigma_n\|_{\mathrm{BV}} = 0$ holds with probability one. Thus, by repeated application of the triangle inequality, $\lim_{n\to\infty}\|\sigma'_n - \sigma^*\|_{\mathrm{BV}} = 0$, i.e.,

$$\mathbb{P}\left(\left\{\lim_{n\to\infty}\|\sigma'_n - \sigma^*\|_{\mathrm{BV}} = 0\right\}\right) = 1.$$

Then, by the robustness of the optimal path σ^*, it follows that

$$\mathbb{P}\left(\left\{\lim_{n\to\infty} c(\sigma'_n) = c^*\right\}\right) = 1.$$

That is the costs of the paths $\{\sigma'_n\}_{n\in\mathbb{N}}$ converges to the optimal cost almost surely, as the number of samples approaches infinity. ☐

D Proof of Theorem 35 (Asymptotic Optimality of k-nearest PRM*)

The proof of this theorem is similar to that of Theorem 34. For the reader's convenience, a complete proof is provided at the expense of repeating some of the arguments.

D.1 Outline of the proof

Let σ^* be a robust optimal path with weak δ-clearance. First, define the sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths as in the proof of Theorem 34.

Second, define a sequence $\{q_n\}_{n\in\mathbb{N}}$ and tile σ_n with a set $B_n = \{B_{n,1}, B_{n,2}, \ldots, B_{n,M}\}$ of overlapping balls of radius q_n. See Figures 23 and 26. Let $x_m \in B_{n,m}$ and $x_{m+1} \in B_{n,m+1}$ be any two points from subsequent balls in B_n. Construct B_n such that the straight path connecting x_m and x_{m+1} lies entirely inside the obstacle free space. Also, construct a set B'_n of balls such that (i) $B'_{n,m}$ and $B_{n,m}$ are centered at the same point and (ii) $B_{n,m}$ contains $B_{n,m}$, and $B_{n,m+1}$, for all $m \in \{1, 2, \ldots, M_n - 1\}$.

Let A_n denote the event that each ball in B_n contains at least one vertex, and A'_n denote the event that each ball in B'_n contains at most $k(n)$ vertices of the graph returned by the k-nearest PRM* algorithm. Third, show that A_n and A'_n occur together for all large n, with probability one. Clearly, this implies that the PRM* algorithm will connect vertices in subsequent ball in B_n with an edge, and any path formed by connecting such vertices will be collision-free.

Finally, show that any sequence of paths formed in this way converges to σ^*. Using the robustness of σ^*, show that the best path in the graph returned by the k-nearest PRM* algorithm converges to $c(\sigma^*)$ almost surely.

D.2 Construction of the sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths

Let $\theta_1, \theta_2 \in \mathbb{R}_{>0}$ be two constants, the precise values of which will be provided shortly. Define

$$\delta_n := \min\left\{\delta,\ (1+\theta_1)\left(\frac{(1+1/d+\theta_2)\,\mu(X_{\mathrm{free}})}{\zeta_d}\right)^{1/d}\left(\frac{\log n}{n}\right)^{1/d}\right\}.$$

Since $\lim_{n\to\infty}\delta_n = 0$ and $0 \le \delta_n \le \delta$ for all $n \in \mathbb{N}$, by Lemma 50, there exists a sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths such that $\lim_{n\to\infty}\sigma_n = \sigma^*$ and σ_n is strongly δ_n-clear for all $n \in \mathbb{N}$.

D.3 Construction of the sequence $\{B_n\}_{n\in\mathbb{N}}$ of sets of balls

Define

$$q_n := \frac{\delta_n}{1+\theta_1}.$$

For each $n \in \mathbb{N}$, use Definition 51 to construct a set $B_n = \{B_{n,1}, B_{n,2}, \ldots, B_{n,M_n}\}$ of overlapping balls that collectively cover σ_n as $B_n := \texttt{CoveringBalls}(\sigma_n, q_n, \theta_1 q_n)$ (see Figures 23 and 26 for an illustration).

D.4 The probability that each ball in B_n contains at least one vertex

Recall that $G_n^{k\mathrm{PRM}^*} = (V_n^{k\mathrm{PRM}^*}, E_n^{k\mathrm{PRM}^*})$ denotes the graph returned by the k-nearest PRM* algorithm, when the algorithm is run with n samples. Let $A_{n,m}$ denote the event that the ball $B_{n,m}$ contains at least one vertex from $V_n^{k\mathrm{PRM}^*}$, i.e., $A_{n,m} = \left\{B_{n,m} \cap V_n^{k\mathrm{PRM}^*} \neq \emptyset\right\}$. Let A_n denote the event that all balls in $B_{n,m}$ contains at least one vertex of $G_n^{k\mathrm{PRM}^*}$, i.e., $A_n = \bigcap_{m=1}^{M_n} A_{n,m}$.

Recall that A_n^c denotes the complement of the event A_n, $\mu(\cdot)$ denotes the Lebesgue measure, and ζ_d is the volume of the unit ball in the d-dimensional Euclidean space. Let s_n denote the length of σ_n.

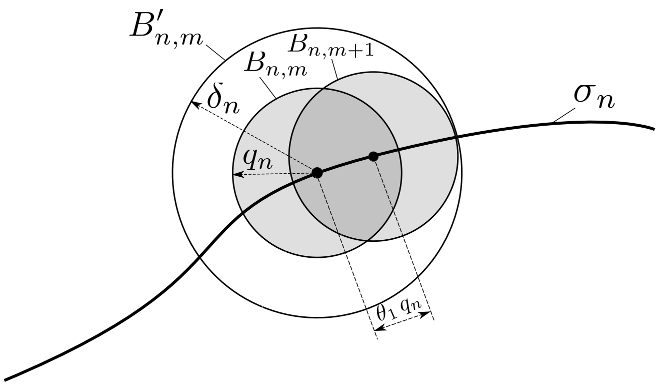

Figure 26: An illustration of the covering balls for the k-nearest PRM* algorithm. The δ_n ball is guaranteed to contain the balls $B_{n,m}$ and $B_{n,m+1}$.

Lemma 57 *For all $\theta_1, \theta_2 > 0$,*

$$\mathbb{P}(A_n^c) \;\leq\; \frac{s_n}{\theta_1} \left(\frac{\zeta_d}{\theta_1 \, (1 + 1/d + \theta_2) \, \mu(X_{\text{free}})} \right)^{1/d} \frac{1}{(\log n)^{1/d} \;\; n^{1+\theta_2}}.$$

In particular, $\sum_{n=1}^{\infty} \mathbb{P}(A_n^c) < \infty$ for all $\theta_1, \theta_2 > 0$.

Proof Let $n_0 \in \mathbb{N}$ be a number for which $\delta_n < \delta$ for all $n > n_0$. A bound on the number of balls in B_n can computed as follows. For all $n > n_0$,

$$M_n = |B_n| \leq \frac{s_n}{\theta_1 \, q_n} = \frac{s_n}{\theta_1} \left(\frac{\zeta_d}{(1 + 1/d + \theta_2) \mu(X_{\text{free}})} \right)^{1/d} \left(\frac{n}{\log n} \right)^{1/d}.$$

The volume of each ball B_n can be computed as

$$\mu(B_{n,m}) = \zeta_d (q_n)^d = (1 + 1/d + \theta_2) \, \mu(X_{\text{free}}) \frac{\log n}{n}.$$

The probability that the ball $B_{n,m}$ does not contain a vertex of the k-nearest PRM* algorithm can be bounded as

$$\mathbb{P}(A_{n,m}^c) = \left(1 - \frac{\mu(B_{n,m})}{\mu(X_{\text{free}})} \right)^n = \left(1 - (1 + 1/d + \theta_2) \frac{\log n}{n} \right)^n \leq n^{-(1 + 1/d + \theta_2)}.$$

Finally, the probability that at least one of the balls in B_n contains no vertex of the k-nearest PRM* can be bounded as

$$\begin{aligned} \mathbb{P}(A_n) \;&=\; \mathbb{P}\left(\bigcup\nolimits_{m=1}^{M_n} A_{n,m} \right) \leq \sum_{m=1}^{M_n} \mathbb{P}(A_{n,m}) = M_n \, \mathbb{P}(A_{n,1}) \\ &\leq\; \frac{s_n}{\theta_1} \left(\frac{\zeta_d}{(1 + 1/d + \theta_2) \, \mu(X_{\text{free}})} \right)^{1/d} \left(\frac{n}{\log n} \right)^{1/d} n^{-(1 + 1/d + \theta_2)} \\ &=\; \frac{s_n}{\theta_1} \left(\frac{\zeta_d}{(1 + 1/d + \theta_2) \, \mu(X_{\text{free}})} \right)^{1/d} \frac{1}{(\log n)^{1/d} \;\; n^{1+\theta_2}}. \end{aligned}$$

Clearly, $\sum_{n=1}^{\infty} \mathbb{P}(A_n^c) < \infty$ for all $\theta_1, \theta_2 > 0$. ☐

D.5 Construction of the sequence $\{B'_n\}_{n \in \mathbb{N}}$ of sets of balls

Construct a set $B'_n = \{B_{n,1}, B_{n,2}, \ldots, B_{n,M_n}\}$ of balls as $B'_n := \texttt{CoveringBalls}(\sigma_n, \delta_n, \theta_1 q_n)$ so that each ball in B'_n has radius δ_n and the spacing between two balls is $\theta_1 q_n$ (see Figure 26).

Clearly, the centers of balls in B'_n coincide with the centers of the balls in B_n, i.e., the center of $B'_{n,m}$ is the same as the center of $B_{n,m}$ for all $m \in \{1, 2, \ldots, M_n\}$ and all $n \in \mathbb{N}$. However, the balls in B'_n have a larger radius than those in B_n.

D.6 The probability that each ball in B'_n contains at most $k(n)$ vertices

Recall that the k-nearest PRM algorithm connects each vertex in the graph with its $k(n)$ nearest vertices when the algorithm is run with n samples, where $k(n) = k_{\mathrm{PRM}} \log n$. Let A'_n denote the event that all balls in B'_n contain at most $k(n)$ vertices of $G_n^{k\mathrm{PRM}^*}$.

Recall that A'^c_n denotes the complement of the event A_n.

Lemma 58 *If $k_{\mathrm{PRM}} > e\,(1 + 1/d)$, then there exists some $\theta_1, \theta_2 > 0$ such that*

$$\mathbb{P}(A'^c_n) \leq \frac{s_n}{\theta_1} \left(\frac{\zeta_d}{(1 + 1/d + \theta_2)\mu(X_{\mathrm{free}})} \right)^{1/d} \frac{1}{(\log n)^{1/d} \;\; n^{-(1+\theta_1)^d(1+1/d+\theta_2)}}.$$

In particular, $\sum_{n=1}^{\infty} \mathbb{P}(A'^c_n) < \infty$ for some $\theta_1, \theta_2 > 0$.

Proof Let $n_0 \in \mathbb{N}$ be a number for which $\delta_n < \delta$ for all $n > n_0$. As shown in the proof of Lemma 57, the number of balls in B'_n satisfies

$$M_n \;=\; |B'_n| \;\leq\; \frac{s_n}{\theta_1 q_n} \;=\; \frac{s_n}{\theta_1} \left(\frac{\zeta_d}{(1 + 1/d + \theta_2)\mu(X_{\mathrm{free}})} \right)^{1/d} \left(\frac{n}{\log n} \right)^{1/d}.$$

For all $n > n_0$, the volume of $B'_{n,m}$ can be computed as

$$\mu(B'_{n,m}) = \zeta_d\,(\delta_n)^d = (1 + \theta_1)^d\,(1 + 1/d + \theta_2)\,\mu(X_{\mathrm{free}})\,\frac{\log n}{n}.$$

Let $I_{n,m,i}$ denote the indicator random variable of the event that sample i falls into ball $B'_{n,m}$. The expected value of $I_{n,m,i}$ can be computed as

$$\mathbb{E}[I_{n,m,i}] = \frac{\mu(B'_{n,m})}{\mu(X_{\mathrm{free}})} = (1 + \theta_1)^d\,(1 + 1/d + \theta_2)\,\frac{\log n}{n}.$$

Let $N_{n,m}$ denote the number of vertices that fall inside the ball $B'_{n,m}$, i.e., $N_{n,m} = \sum_{i=1}^{n} I_{n,m,i}$. Then,

$$\mathbb{E}[N_{n,m}] = \sum_{i=1}^{n} \mathbb{E}[I_{n,m,i}] = n\;\mathbb{E}[I_{n,m,1}] = (1 + \theta_1)^d(1 + 1/d + \theta_2)\log n.$$

Since $\{I_{n,m,i}\}_{i=1}^{n}$ are independent identically distributed random variables, large deviations of their sum, $M_{n,m}$, can be bounded by the following Chernoff bound (Dubhashi and Panconesi, 2009):

$$\mathbb{P}\big(\,\{\, N_{n,m} > (1 + \epsilon)\,\mathbb{E}[N_{n,m}] \,\}\,\big) \;\leq\; \left(\frac{e^{\epsilon}}{(1 + \epsilon)^{(1+\epsilon)}} \right)^{\mathbb{E}[N_{n,m}]},$$

for all $\epsilon > 0$. In particular, for $\epsilon = e - 1$,

$$\mathbb{P}\big(\,\{\, N_{n,m} > e\,\mathbb{E}[N_{n,m}] \,\}\,\big) \;\leq\; e^{-\mathbb{E}[N_{n,m}]} \;=\; e^{-(1+\theta_1)^d(1+1/d+\theta_2)\log n} \;=\; n^{-(1+\theta_1)^d(1+1/d+\theta_2)}.$$

Since $k(n) > e\,(1{+}1/d)\,\log n$, there exists some $\theta_1, \theta_2 > 0$ independent of n such that $e\,\mathbb{E}[N_{n,k}] = e\,(1+\theta_1)\,(1+1/d+\theta_2)\,\log n \leq k(n)$. Then, for the same values of θ_1 and θ_2,

$$\mathbb{P}\big(\,\{\,N_{n,m} > k(n)\,\}\,\big) \;\leq\; \mathbb{P}\big(\,\{\,N_{n,m} > e\,\mathbb{E}[N_{n,m}]\,\}\,\big) \;\leq\; n^{-(1+\theta_1)^d(1+1/d+\theta_2)}.$$

Finally, consider the probability of the event that at least one ball in B_n contains more than $k(n)$ nodes. Using the union bound together with the inequality above

$$\mathbb{P}\left(\bigcup_{m=1}^{M_n}\{N_{n,m} > k(n)\}\right) \;\leq\; \sum_{m=1}^{M_n}\mathbb{P}\big(\{N_{n,m} > k(n)\}\big) \;=\; M_n\;\mathbb{P}\big(\{N_{n,1} > k(n)\}\big)$$

Hence,

$$\mathbb{P}(A_n'^c) = \mathbb{P}\left(\bigcup_{m=1}^{M_n}\{N_{n,m} > k(n)\}\right) \;\leq\; \frac{s_n}{\theta_1}\left(\frac{\zeta_d}{(1+1/d+\theta_2)\mu(X_{\text{free}})}\right)^{1/d}\frac{1}{(\log n)^{1/d}\;\;n^{-(1+\theta_1)^d(1+1/d+\theta_2)}}.$$

Clearly, $\sum_{n=1}^{\infty}\mathbb{P}(A_n'^c) < \infty$ for the same values of θ_1 and θ_2. □

D.7 Connecting the vertices in the subsequent balls in B_n

First, note the following lemma.

Lemma 59 *If $k_{\text{PRM}} > e\,(1+1/d)^{1/d}$, then there exists $\theta_1, \theta_2 > 0$ such that the event that each ball in B_n contains at least one vertex and each ball in B_n' contains at most $k(n)$ vertices occurs for all large n, with probability one, i.e.,*

$$\mathbb{P}\left(\liminf_{n\to\infty}(A_n \cap A_n')\right) = 1.$$

Proof Consider the event $A_n^c \cup A_n'^c$, which is the complement of $A_n \cap A_n'$. Using the union bound,

$$\mathbb{P}\left(A_n^c \cup A_n'^c\right) \leq \mathbb{P}(A_n^c) + \mathbb{P}(A_n'^c).$$

Summing both sides,

$$\sum_{n=1}^{\infty}\mathbb{P}(A_n^c \cup A_N'^c) \leq \sum_{n=1}^{\infty}\mathbb{P}(A_n^c) + \sum_{n=1}^{\infty}\mathbb{P}(A_n'^c) < \infty,$$

where the last inequality follows from Lemmas 57 and 58. Then, by the Borel-Cantelli lemma, $\mathbb{P}\left(\limsup_{n\to\infty}(A_n^c \cup A_n'^c)\right) = \mathbb{P}\left(\limsup_{n\to\infty}(A_n \cap A_n')^c\right) = 0$, which implies $\mathbb{P}\left(\liminf_{n\to\infty}(A_n \cap A_n')\right) = 1$. □

Note that for each $m \in \{1, 2, \ldots, M_n - 1\}$, both $B_{n,m}$ and $B_{n,m+1}$ lies entirely inside the ball $B_{n,m}'$ (see Figure 26). Hence, whenever the balls $B_{n,m}$ and $B_{n,m+1}$ contain at least one vertex each, and $B_{n,m}'$ contains at most $k(n)$ vertices, the k-nearest PRM* algorithm attempts to connect all vertices in $B_{n,m}$ and $B_{n,m+1}$ with one another.

The following lemma guarantees that connecting any two points from two consecutive balls in B_n results in a collision-free trajectory. The proof of the lemma is essentially the same as that of Lemma 54.

Lemma 60 *For all $n \in \mathbb{N}$ and all $m \in \{1, 2, \ldots, M_n\}$, if $x_m \in B_{n,m}$ and $x_{m+1} \in B_{n,m+1}$, then the line segment connecting x_m and x_{m+1} lies in the obstacle-free space, i.e.,*

$$\alpha\,x_m + (1-\alpha)\,x_{m+1} \in X_{\text{free}}, \qquad \textit{for all } \alpha \in [0,1].$$

D.8 Convergence to the optimal path

The proof of the following lemma is similar to that of Lemma 55, and is omitted here.

Let P_n denote the set of all paths in the graph returned by k-PRM* algorithm at the end of n iterations. Let σ_n' be the path that is closest to σ_n in terms of the bounded variation norm among all those paths in P_n, i.e., $\sigma_n' := \min_{\sigma' \in P_n} \|\sigma' - \sigma_n\|$.

Lemma 61 *The random variable* $\|\sigma_n' - \sigma_n\|_{\mathrm{BV}}$ *converges to zero almost surely, i.e.,*

$$\mathbb{P}\left(\left\{\lim_{n\to\infty} \|\sigma_n' - \sigma_n\|_{\mathrm{BV}} = 0\right\}\right) = 1.$$

A corollary of the lemma above is that $\lim_{n\to\infty} \sigma_n' = \sigma^*$ with probability one. Then, the result follows by the robustness of the optimal solution (see the proof of Lemma 56 for details).

E Proof of Theorem 36 (Asymptotic optimality of RRG)

E.1 Outline of the proof

The proof of this theorem is similar to that of Theorem 34. The main difference is the definition of C_n that denotes the event that the RRG algorithm has sufficiently explored the obstacle free space. More precisely, C_n is the event that for any point x in the obstacle free space, the graph maintained by the RRG algorithm algorithm includes a vertex that can be connected to x.

Construct the sequence $\{\sigma_n\}_{n\in\mathbb{N}}$ of paths and the sequence $\{B_n\}_{n\in\mathbb{N}}$ of balls as in the proof of Theorem 34. Let A_n denote the event that each ball in B_n contains a vertex of the graph maintained by the RRG by the end of iteration n. Compute n by conditioning on the event that C_i holds for all $i \in \{\lfloor \theta_3\, n \rfloor, \ldots, n\}$, where $0 < \theta_3 < 1$ is a constant. Show that the probability that C_i fails to occur for any such i is small enough to guarantee that A_n occurs for all large n with probability one. Complete the proof as in the proof of Theorem 34.

E.2 Definitions of $\{\sigma_n\}_{n\in\mathbb{N}}$ and $\{B_n\}_{n\in\mathbb{N}}$

Let $\theta_1 > 0$ be a constant. Define δ_n, σ_n, q_n, and B_n as in the proof of Theorem 34.

E.3 Probability that each ball in B_n contains at least one vertex

Let $A_{n,m}$ be the event that the ball $B_{n,m}$ contains at least one vertex of the RRG at the end of n iterations. Let A_n be the event that all balls in B_n contain at least one vertex of the RRG at the end of iteration n, i.e., $A_n = \bigcap_{m=1}^{M_n} A_{n,m}$, where M_n is the number of balls in B_n. Recall that γ_{RRG} is the constant used in defining the connection radius of the RRG algorithm (see Algorithm 5).

Lemma 62 *If* $\gamma_{\mathrm{RRG}} > 2(1+1/d)^{1/d}\left(\frac{\mu(X_{\mathrm{free}})}{\zeta_d}\right)^{1/d}$ *then there exists* $\theta_1 > 0$ *such that* A_n *occurs for all large* n *with probability one, i.e.,*

$$\mathbb{P}\left(\liminf_{n\to\infty} A_n\right) = 1.$$

The proof of this lemma requires two intermediate results, which are provided next.

Recall that η is the parameter used in the `Steer` procedure (see the definition of `Steer` procedure in Section 3.1). Let C_n denote the event that for any point $x \in X_{\mathrm{free}}$, the graph returned by the RRG algorithm includes a vertex v such that $\|x - v\| \leq \eta$ and the line segment joining v and x is collision-free. The following lemma establishes an bound on the probability that this event fails to occur at iteration n.

Lemma 63 *There exists constants* $a, b \in \mathbb{R}_{>0}$ *such that* $P(C_n^c) \leq a\, e^{-b\,n}$ *for all* $n \in \mathbb{N}$.

Proof Partition $X_{\rm free}$ into finitely many convex sets such that each partition is bounded by a ball a radius η. Such a finite partition exists by the boundedness of $X_{\rm free}$. Denote this partition by $X_1', X_2', \dots, X_M'$. Since the probability of failure decays to zero with an exponential rate, for any $m \in \{1, 2, \dots, M\}$, the probability that X_m' fails to contain a vertex of the RRG decays to zero with an exponential rate, i.e.,

$$\mathbb{P}\left(\{\nexists x \in V_n^{\rm RRG} \cap X_m'\}\right) \leq a_m\, e^{-b_m\, n}$$

The probability that at least one partition fails to contain one vertex of the RRG also decays to zero with an exponential rate. That is, there exists $a, b \in \mathbb{R}_{>0}$ such that

$$\mathbb{P}\left(\bigcup\nolimits_{m=1}^{M} \{\nexists x \in V_n^{\rm RRG} \cap X_m'\}\right) \leq \sum_{m=1}^{M} \mathbb{P}\left(\{\nexists x \in V_n^{\rm RRG} \cap X_m'\}\right) \leq \sum_{m=1}^{M} a_m\, e^{-b_m n} \leq a\, e^{-b\, n},$$

where the first inequality follows from the union bound. $\square$

Let $0 < \theta_3 < 1$ be a constant independent of n. Consider the event that C_i occurs for all i that is greater than $\theta_3\, n$, i.e., $\bigcap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i$. The following lemma analyzes the probability of the event that $\bigcap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i$ fails to occur.

Lemma 64 *For any* $\theta_3 \in (0, 1)$*,*

$$\sum_{n=1}^{\infty} \mathbb{P}\left(\left(\bigcap\nolimits_{i=\lfloor \theta_3 n \rfloor}^{n} C_i\right)^c\right) \;<\; \infty.$$

Proof The following inequalities hold:

$$\sum_{n=1}^{\infty} \mathbb{P}\left(\left(\bigcap\nolimits_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right)^c\right) \;=\; \sum_{n=1}^{\infty} \mathbb{P}\left(\bigcup\nolimits_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i^c\right) \;\leq\; \sum_{n=1}^{\infty} \sum_{i=\lfloor \theta_3\, i \rfloor}^{n} \mathbb{P}(C_i^c) \;\leq\; \sum_{n=1}^{\infty} \sum_{i=\lfloor \theta_3\, n \rfloor}^{n} a\, e^{-b\, i},$$

where the last inequality follows from Lemma 63. The right-hand side is finite for all $a, b > 0$. $\square$

Proof of Lemma 62 It is shown that $\sum_{n=1}^{\infty} \mathbb{P}\left(A_n^c\right) < \infty$, which, by the Borel-Cantelli Lemma (Grimmett and Stirzaker, 2001), implies that A_n^c occurs infinitely often with probability zero, i.e., $\mathbb{P}(\limsup_{n\to\infty} A_n^c) = 0$, which in turn implies $\mathbb{P}(\liminf_{n\to\infty} A_n) = 1$.

Let $n_0 \in \mathbb{N}$ be a number for which $\delta_n < \delta$ for all $n > n_0$. First, for all $n > n_0$, the number of balls in B_n can be bounded by (see the proof of Lemma 52 for details)

$$M_n = |B_n| \leq \frac{(2 + \theta_1)\, s_n}{\theta_1\, \gamma_{\rm RRG}} \left(\frac{n}{\log n}\right)^{1/d}.$$

Second, for all $n > n_0$, the volume of each ball in B_n can be calculated as (see the proof of Lemma 52)

$$\mu(B_{n,m}) \;=\; \zeta_d \left(\frac{\gamma_{\rm PRM}}{2 + \theta_1}\right)^d \frac{\log n}{n},$$

where ζ_d is the volume of the unit ball in the d-dimensional Euclidean space.

Third, conditioning on the event $\bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i$, each new sample will be added to the graph maintained by the RRG algorithm as a new vertex between iterations $i = \lfloor \theta_3\, n \rfloor$ and $i = n$. Thus,

$$\begin{aligned}
\mathbb{P}\left(A_{n,m}^c \,\Big|\, \bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i\right) &\le \left(1 - \frac{\mu(B_{n,m})}{\mu(X_{\mathrm{free}})}\right)^{n-\lfloor\theta_3\, n\rfloor} \le \left(1 - \frac{\mu(B_{n,m})}{\mu(X_{\mathrm{free}})}\right)^{(1-\theta_3)\, n} \\
&\le \left(1 - \frac{\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{RRG}}}{2+\theta_1}\right)^d \frac{\log n}{n}\right)^{(1-\theta_3)\, n} \\
&\le e^{-\frac{(1-\theta_3)\,\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{RRG}}}{2+\theta_1}\right)^d \log n} \le n^{-\frac{(1-\theta_3)\,\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{RRG}}}{2+\theta_1}\right)^d},
\end{aligned}$$

where the fourth inequality follows from $(1-1/f(n))^{g(n)} \le e^{g(n)/f(n)}$.

Fourth,

$$\begin{aligned}
\mathbb{P}\left(A_n^c \,\Big|\, \bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i\right) &\le \mathbb{P}\left(\bigcup_{m=1}^{M_n} A_{n,m}^c \,\Big|\, \bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i\right) \\
&\le \sum_{m=1}^{M_n} \mathbb{P}\left(A_{n,m}^c \,|\, \bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i\right) \\
&= M_n\; \mathbb{P}\left(A_{n,1}^c \,|\, \bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i\right) \\
&\le \frac{(2+\theta_1)\, s_n}{\theta_1\, \gamma_{\mathrm{RRG}}}\left(\frac{n}{\log n}\right)^{1/d} n^{-\frac{(1-\theta_3)\,\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{RRG}}}{2+\theta_1}\right)^d}.
\end{aligned}$$

Hence,

$$\sum_{n=1}^{\infty} \mathbb{P}\left(A_n^c \,\Big|\, \bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i\right) < \infty,$$

whenever $\frac{(1-\theta_3)\,\zeta_d}{\mu(X_{\mathrm{free}})}\left(\frac{\gamma_{\mathrm{RRG}}}{2+\theta_1}\right)^d - 1/d > 1$, i.e., $\gamma_{\mathrm{RRG}} > (2+\theta_1)(1+1/d)^{1/d}\left(\frac{\mu(X_{\mathrm{free}})}{(1-\theta_3)\,\zeta_d}\right)^{1/d}$, which is satisfied by appropriately choosing the constants θ_1 and θ_3, since $\gamma_{\mathrm{RRG}} > 2\,(1+1/d)^{1/d}\left(\frac{\mu(X_{\mathrm{free}})}{\zeta_d}\right)^{1/d}$.

Finally,

$$\begin{aligned}
\mathbb{P}\left(A_n^c \,|\, \bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i\right) &= \frac{\mathbb{P}\left(A_n^c \cap \left(\cap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i\right)\right)}{\mathbb{P}\left(\bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i\right)} \\
&\ge \mathbb{P}(A_n^c \cap (\cap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i)) \\
&= 1 - \mathbb{P}(A_n \cup (\cap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i)^c) \\
&\ge 1 - \mathbb{P}(A_n) - \mathbb{P}\big((\cap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i)^c\big) \\
&= \mathbb{P}(A_n^c) - \mathbb{P}\big((\cap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_i)^c\big).
\end{aligned}$$

Taking the infinite sum of both sides yields

$$\sum_{n=1}^{\infty} \mathbb{P}(A_n^c) \le \sum_{n=1}^{\infty} \mathbb{P}\left(A_n^c \,|\, \bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_n\right) + \sum_{n=1}^{\infty} \mathbb{P}\left(\left(\bigcap_{i=\lfloor\theta_3\, n\rfloor}^{n} C_n\right)^c\right).$$

The first term on the right hand side is shown to be finite above. The second term is finite by Lemma 64. Hence, $\sum_{n=1}^{\infty} \mathbb{P}(A_n) < \infty$. Then, by the Borel Cantelli lemma, A_n^c occurs infinitely often with probability zero, which implies that its complement A_n occurs for all large n, with probability one. □

E.4 Convergence to the optimal path

The proof of the following lemma is similar to that of Lemma 55, and is omitted here.

Let P_n denote the set of all paths in the graph returned by RRG algorithm at the end of n iterations. Let σ'_n be the path that is closest to σ_n in terms of the bounded variation norm among all those paths in P_n, i.e., $\sigma'_n := \min_{\sigma' \in P_n} \|\sigma' - \sigma_n\|$.

Lemma 65 *The random variable $\|\sigma'_n - \sigma_n\|_{\mathrm{BV}}$ converges to zero almost surely, i.e.,*

$$\mathbb{P}\left(\left\{\lim_{n\to\infty} \|\sigma'_n - \sigma_n\|_{\mathrm{BV}} = 0\right\}\right) = 1.$$

A corollary of the lemma above is that $\lim_{n\to\infty} \sigma'_n = \sigma^*$ with probability one. Then, the result follows by the robustness of the optimal solution (see the proof of Lemma 56 for details).

F Proof of Theorem 37 (asymptotic optimality of k-nearest RRG)

F.1 Outline of the proof

The proof of this theorem is a combination of that of Theorem 35 and 36.

Define the sequences $\{\sigma_n\}_{n\in\mathbb{N}}$, $\{B_n\}_{n\in\mathbb{N}}$, and $\{B'_n\}_{n\in\mathbb{N}}$ as in the proof of Theorem 35. Define the event C_n as in the proof of Theorem 36. Let A_n denote the event that each ball in B_n contains at least one vertex, and A'_n denote the event that each ball in B'_n contains at most $k(n)$ vertices of the graph maintained by the RRG algorithm, by the end of iteration n. Compute A_n and A'_n by conditioning on the event that C_i holds for all $i = \theta_3\, n$ to n. Show that this is enough to guarantee that A_n and A'_n hold together for all large n, with probability one.

F.2 Definitions of $\{\sigma_n\}_{n\in\mathbb{N}}$, $\{B_n\}_{n\in\mathbb{N}}$, and $\{B'_n\}_{n\in\mathbb{N}}$

Let $\theta_1, \theta_2 > 0$ be two constants. Define δ_n, σ_n, q_n, B_n, and B'_n as in the proof of Theorem 35.

F.3 The probability that each ball in B_n contains at least one vertex

Let $A_{n,m}$ denote the event that the ball $B_{n,m}$ contains at least one vertex of the graph maintained by the k-nearest RRG algorithm by the end of iteration n. Let A_n denote the event that all balls in $B_{n,m}$ contain at least one vertex of the same graph, i.e., $A_n = \bigcup_{m=1}^{M_n} A_{n,m}$. Let s_n denote the length of σ_n, i.e., $TV(\sigma_n)$. Recall η is the parameter in the `Steer` procedure. Let C_n denote the event that for any point $x \in X_{\mathrm{free}}$, the k-nearest RRG algorithm includes a vertex v such that $\|x - v\| \leq \eta$.

Lemma 66 *For any $\theta_1, \theta_2 > 0$ and any $\theta_3 \in (0,1)$,*

$$\mathbb{P}\left(A_n^c \,\Big|\, \bigcap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right) \;\leq\; \frac{s_n}{\theta_1}\left(\frac{\zeta_d}{\theta_1\,(1 + 1/d + \theta_2)\,\mu(X_{\mathrm{free}})}\right)^{1/d} \frac{1}{(\log n)^{1/d}\;\; n^{(1-\theta_3)(1+1/d+\theta_2)-1/d}}.$$

In particular, $\sum_{n=1}^{\infty} \mathbb{P}(A_n^c \mid \bigcap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i) < \infty$ for any $\theta_1, \theta_2 > 0$ and some $\theta_3 \in (0,1)$.

Proof Let $n_0 \in \mathbb{N}$ be a number for which $\delta_n < \delta$ for all $n > n_0$. Then, for all $n > n_0$,

$$M_n = |B_n| \leq \frac{s_n}{\theta_1\, q_n} = \frac{s_n}{\theta_1}\left(\frac{\zeta_d}{(1 + 1/d + \theta_2)\mu(X_{\mathrm{free}})}\right)^{1/d}\left(\frac{n}{\log n}\right)^{1/d}.$$

The volume of each ball B_n can be computed as

$$\mu(B_{n,m}) = \zeta_d (q_n)^d = (1 + 1/d + \theta_2)\, \mu(X_{\mathrm{free}}) \frac{\log n}{n}.$$

Given $\bigcap_{i=\lceil \theta_3\, n \rceil}^{n} C_i$, the probability that the ball $B_{n,m}$ does not contain a vertex of the k-nearest PRM* algorithm can be bounded as

$$\mathbb{P}\Big(A^c_{n,m} \,\big|\, \bigcap\nolimits_{i=\lceil \theta_3\, n \rceil}^{n} C_i\Big) \;=\; \left(1 - \frac{\mu(B_{n,m})}{\mu(X_{\mathrm{free}})}\right)^{(1-\theta_3)\, n} \;=\; \left(1 - (1+1/d+\theta_2) \frac{\log n}{n}\right)^{(1-\theta_3)\, n} \;\leq n^{-(1-\theta_3)\,(1+1/d+\theta_2)}.$$

Finally, the probability that at least one of the balls in B_n contains no vertex of the k-nearest PRM* can be bounded as

$$\begin{aligned}
\mathbb{P}(A_n) \;&=\; \mathbb{P}\left(\bigcup\nolimits_{m=1}^{M_n} A_{n,m}\right) \;\leq\; \sum_{m=1}^{M_n} \mathbb{P}(A_{n,m}) = M_n\, \mathbb{P}(A_{n,1}) \\
&\leq\; \frac{s_n}{\theta_1} \left(\frac{\zeta_d}{(1+1/d+\theta_2)\, \mu(X_{\mathrm{free}})}\right)^{1/d} \frac{1}{(\log n)^{1/d}\;\; n^{(1-\theta_3)\,(1+1/d+\theta_2)-1/d}}.
\end{aligned}$$

Clearly, for all $\theta_1, \theta_2 > 0$, there exists some $\theta_3 \in (0,1)$ such that $\sum_{n=1}^{\infty} \mathbb{P}(A^c_n) < \infty$. ☐

F.4 The probability that each ball in B'_n contains at most $k(n)$ vertices

Let A'_n denote the event that all balls in B'_n contain at most $k(n)$ vertices of the graph maintained by the RRG algorithm, by end of iteration n.

Lemma 67 *If $k_{\mathrm{PRM}} > e\,(1+1/d)$, then there exists $\theta_1, \theta_2, \theta_3 > 0$ such that*

$$\mathbb{P}\left(A'^c_n \,\Big|\, \bigcap\nolimits_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right) \;\leq\; \frac{s_n}{\theta_1} \left(\frac{\zeta_d}{(1+1/d+\theta_2)\mu(X_{\mathrm{free}})}\right)^{1/d} \frac{1}{(\log n)^{1/d}\;\; n^{-(1-\theta_3)(1+\theta_1)^d(1+1/d+\theta_2)}}.$$

In particular, $\sum_{n=1}^{\infty} \mathbb{P}(A^c_n \mid \bigcap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i) < \infty$ for some $\theta_1, \theta_2 > 0$ and some $\theta_3 > 0$.

Proof Let $n_0 \in \mathbb{N}$ be a number for which $\lambda_n < \delta$ for all $n > n_0$. Then, the number of balls in B'_n and the volume of each ball can be computed as

$$M_n \;=\; |B'_n| \;\leq\; \frac{s_n}{\theta_1 q_n} \;=\; \frac{s_n}{\theta_1} \left(\frac{\zeta_d}{(1+1/d+\theta_2)\mu(X_{\mathrm{free}})}\right)^{1/d} \left(\frac{n}{\log n}\right)^{1/d}.$$

$$\mu(B'_{n,m}) = \zeta_d\, (\lambda_n)^d = (1+\theta_1)^d\, (1+1/d+\theta_2)\, \mu(X_{\mathrm{free}})\, \frac{\log n}{n}.$$

Let $I_{n,m,i}$ denote the indicator random variable of the event that sample i falls into ball $B'_{n,m}$. The expected value of $I_{n,m,i}$ can be computed as

$$\mathbb{E}[I_{n,m,i}] = \frac{\mu(B'_{n,m})}{\mu(X_{\mathrm{free}})} = (1+\theta_1)^d\, (1+1/d+\theta_2)\, \frac{\log n}{n}.$$

Let $N_{n,m}$ denote the number of vertices that fall inside the ball $B'_{n,m}$ between iterations $\lfloor \theta_3\, n \rfloor$ and n, i.e., $N_{n,m} = \sum_{i=\lfloor \theta_3\, n \rfloor}^{n} I_{n,m,i}$. Then,

$$\mathbb{E}[N_{n,m}] = \sum_{i=\lfloor \theta_3\, n \rfloor}^{n} \mathbb{E}[I_{n,m,i}] = (1-\theta_3)\, n\; \mathbb{E}[I_{n,m,1}] = (1-\theta_3)\, (1+\theta_1)^d\, (1+1/d+\theta_2) \log n.$$

Since $\{I_{n,m,i}\}_{i=1}^{n}$ are independent identically distributed random variables, large deviations of their sum, $M_{n,m}$, can be bounded by the following Chernoff bound (Dubhashi and Panconesi, 2009):

$$\mathbb{P}\big(\{\, N_{n,m} > (1+\epsilon)\,\mathbb{E}[N_{n,m}]\,\}\big) \;\leq\; \left(\frac{e^{\epsilon}}{(1+\epsilon)^{(1+\epsilon)}}\right)^{\mathbb{E}[N_{n,m}]},$$

for all $\epsilon > 0$. In particular, for $\epsilon = e - 1$,

$$\mathbb{P}\big(\{\, N_{n,m} > e\,\mathbb{E}[N_{n,m}]\,\}\big) \;\leq\; e^{-\mathbb{E}[N_{n,m}]} \;=\; n^{-(1-\theta_3)\,(1+\theta_1)^d\,(1+1/d+\theta_2)}.$$

Since $k(n) > e\,(1+1/d)\,\log n$, there exists some $\theta_1, \theta_2 > 0$ and $\theta_3 \in (0,1)$, independent of n, such that $e\,\mathbb{E}[N_{n,k}] = e\,(1-\theta_3)\,(1+\theta_1)\,(1+1/d+\theta_2)\,\log n \leq k(n)$. Then, for the same values of θ_1 and θ_2,

$$\mathbb{P}\big(\{\, N_{n,m} > k(n)\,\}\big) \;\leq\; \mathbb{P}\big(\{\, N_{n,m} > e\,\mathbb{E}[N_{n,m}]\,\}\big) \;\leq\; n^{-(1-\theta_3)\,(1+\theta_1)^d\,(1+1/d+\theta_2)}.$$

Finally, consider the probability of the event that at least one ball in B_n contains more than $k(n)$ nodes. Using the union bound together with the inequality above

$$\mathbb{P}\left(\bigcup\nolimits_{m=1}^{M_n}\{N_{n,m} > k(n)\}\right) \;\leq\; \sum_{m=1}^{M_n}\mathbb{P}\big(\{N_{n,m} > k(n)\}\big) \;=\; M_n\,\mathbb{P}\big(\{N_{n,1} > k(n)\}\big)$$

Hence,

$$\begin{aligned}\mathbb{P}\Big(A_n'^{c} \mid \bigcap\nolimits_{i=\lfloor\theta_3\,n\rfloor}^{n} C_i\Big) &= \mathbb{P}\left(\bigcup\nolimits_{m=1}^{M_n}\{N_{n,m} > k(n)\}\right)\\ &\leq \frac{s_n}{\theta_1}\left(\frac{\zeta_d}{(1+1/d+\theta_2)\mu(X_{\mathrm{free}})}\right)^{1/d}\frac{1}{(\log n)^{1/d}\;\;n^{-(1-\theta_3)\,(1+\theta_1)^d(1+1/d+\theta_2)}}.\end{aligned}$$

Clearly, $\sum_{n=1}^{\infty}\mathbb{P}\big(A_n'^{c} \mid \cap_{i=\lfloor\theta_3\,n\rfloor}^{n} C_i\big) < \infty$ for the same values of θ_1, θ_2, and θ_3. □.

F.5 Connecting the vertices in subsequent balls in B_n

Lemma 68 *If $k_{\mathrm{PRM}} > e\,(1+1/d)^{1/d}$, then there exists $\theta_1, \theta_2 > 0$ such that the event that each ball in B_n contains at least one vertex and each ball in B_n' contains at most $k(n)$ vertices occurs for all large n, with probability one, i.e.,*

$$\mathbb{P}\left(\liminf_{n\to\infty}(A_n \cap A_n')\right) = 1.$$

First note the following lemma.

Lemma 69 *For any $\theta_3 \in (0,1)$,*

$$\sum_{n=1}^{\infty}\mathbb{P}\left(\left(\bigcap\nolimits_{i=\lfloor\theta_3 n\rfloor}^{n} C_n\right)^{c}\right) \;<\; \infty.$$

Proof Since the RRG algorithm and the k-nearest RRG algorithm have the same vertex sets, i.e., $V_n^{\mathrm{RRG}} = V_n^{k\mathrm{RRG}}$ surely for all $n \in \mathbb{N}$, the lemma follows from Lemma 64. □

Proof of Lemma 68 Note that

$$\begin{aligned}\mathbb{P}\left((A_n^c \cup A_n'^c) \mid \bigcap\nolimits_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right) \;&=\; \frac{\mathbb{P}\left(A_n^c \cap \left(\cap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right)\right)}{\mathbb{P}\left(\bigcap\nolimits_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right)} \\ &\geq\; \mathbb{P}\left((A_n^c \cup A_n'^c) \cap \left(\cap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right)\right) \\ &\geq\; \mathbb{P}(A_n^c \cup A_n'^c) - \mathbb{P}\big((\cap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i)^c\big),\end{aligned}$$

where the last inequality follows from the union bound. Rearranging and using the union bound,

$$\mathbb{P}(A_n^c \cup A_n'^c) \;\leq\; \mathbb{P}\big(A_n^c \mid \cap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\big) + \mathbb{P}\big(A_n^c \mid \cap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\big) + \mathbb{P}\big(\big(\cap_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\big)^c\big).$$

Summing both sides,

$$\sum_{n=1}^{\infty} \mathbb{P}(A_n^c \cup A_n'^c) \;\leq\; \sum_{n=1}^{\infty} \mathbb{P}\left(A_n^c \mid \bigcap\nolimits_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right) + \sum_{n=1}^{\infty} \mathbb{P}\left(A_n'^c \mid \bigcap\nolimits_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right) + \sum_{n=1}^{\infty} \mathbb{P}\left(\left(\bigcap\nolimits_{i=\lfloor \theta_3\, n \rfloor}^{n} C_i\right)^c\right),$$

where the right hand side is finite by Lemmas 66, 67, and 69, by picking θ_3 close to one. Hence, $\sum_{n=1}^{\infty} \mathbb{P}(A_n^c \cup A_n'^c) < \infty$. Then, by the Borel-Cantelli lemma, $\mathbb{P}(\limsup_{n\to\infty}(A_n^c \cup A_n'^c)) = 0$, or equivalently $\mathbb{P}(\liminf_{n\to\infty}(A_n \cap A_n')) = 1$. □

F.6 Convergence to the optimal path

The proof of the following two lemmas are essentially the same as that of Lemma 55, and is omitted here. Let P_n denote the set of all paths in the graph returned by k-RRG algorithm at the end of n iterations. Let σ_n' be the path that is closest to σ_n in terms of the bounded variation norm among all those paths in P_n, i.e., $\sigma_n' := \min_{\sigma' \in P_n} \|\sigma' - \sigma_n\|$.

Lemma 70 *The random variable* $\|\sigma_n' - \sigma_n\|_{\mathrm{BV}}$ *converges to zero almost surely, i.e.,*

$$\mathbb{P}\left(\left\{\lim_{n\to\infty} \|\sigma_n' - \sigma_n\|_{\mathrm{BV}} = 0\right\}\right) = 1.$$

A corollary of the lemma above is that $\lim_{n\to\infty} \sigma_n' = \sigma^*$ with probability one. Then, the result follows by the robustness of the optimal solution (see the proof of Lemma 56 for details).

G Proof of Theorem 38 (Asymptotic optimality of RRT*)

For simplicity, the proof will assume the steering parameter η to be large enough, i.e., $\eta \geq \mathrm{diam}(\mathcal{X})$, although the results hold for any $\eta > 0$.

G.1 Marked point process

Consider the following marked point process. Let $\{X_1, X_2, \ldots, X_n\}$ be a independent uniformly distributed points drawn from X_{free} and let $\{Y_1, Y_2, \ldots, Y_n\}$ be independent uniform random variables with support $[0, 1]$. Each point X_i is associated with a mark Y_i that describes the order of X_i in the process. More precisely, a point X_i is assumed to be drawn after another point $X_{i'}$ if $Y_{i'} < Y_i$. We will also assume that the point process includes the point x_{init} with mark $Y = 0$.

Consider the graph formed by adding an edge $(X_{i'}, X_i)$, whenever (i) $Y_{i'} < Y_i$ and (ii) $\|X_i - X_{i'}\| \leq r_n$ both hold. Notice that, formed in this way, G_n includes no directed cycles. Denote this graph by $G_n = (V_n, E_n)$. Also, consider a subgraph G'_n of G_n formed as follows. Let $c(X_i)$ denote the cost of best path starting from x_{init} and reaching X_i. In G'_n, each vertex X_i has a single parent X_i with the smallest cost $c(X_i)$. Since the graph is built incrementally, the cost of the best path reaching X_i will be the same as the one reaching $X_{i'}$ in both G_n and G'_n. Clearly, G'_n is equivalent to the graph returned by the RRT* algorithm at the end of n iterations, if the steering parameter η is large enough.

Let Y_n and the Y'_n denote the costs of the best paths starting from x_{init} and reaching the goal region in G_n and G'_n, respectively. Then, $\limsup_{n\to\infty} Y_n = \limsup_{n\to\infty} Y'_n$ surely. In the rest of the proof, it is shown that $\mathbb{P}(\{\limsup_{n\to\infty} Y_n\}) = 1$, which implies that $\mathbb{P}(\{\limsup_{n\to\infty} Y'_n\}) = 1$, which in turn implies the result.

G.2 Definitions of $\{\sigma_n\}_{n\in\mathbb{N}}$ and $\{B_n\}_{n\in\mathbb{N}}$

Let σ^* denote an optimal path. Define

$$\delta_n := \min\{\delta, 4\, r_n\},$$

where r_n is the connection radius of the RRT* algorithm. Let $\{\sigma_n\}_{n\in\mathbb{N}}$ be the sequence paths, the existence of which is guaranteed by Lemma 50.

For each $n \in \mathbb{N}$, construct a sequence $\{B_n\}_{n\in\mathbb{N}}$ of balls that cover σ_n as $B_n = \{B_{n,1}, B_{n,2}, \ldots, B_{n,M_n}\} :=$ `CoveringBalls`$(\sigma_n, r_n, 2\, r_n)$ (see Definition 51), where r_n is the connection radius of the RRT* algorithm, i.e., $r_n = \gamma_{\mathrm{RRT}^*} \left(\frac{\log n}{n}\right)^{1/d}$. Clearly, the balls in B_n are openly disjoint, since the spacing between any two consecutive balls is $2\, r_n$.

G.3 Connecting the vertices in subsequent balls in B_n

For all $m \in \{1, 2, \ldots, M_n\}$, let $A_{n,m}$ denote the event that there exists two vertices $X_i, X_{i'} \in V_n^{\mathrm{RRT}^*}$ such that $X_i \in B_{n,m}$, $X_{i'} \in B_{n,m+1}$ and $Y_{i'} \le Y_i$, where Y_i and $Y_{i'}$ are the marks associated with points X_i and $X_{i'}$, respectively. Notice that, in this case, X_i and $X_{i'}$ will be connected with an edge in G_n. Let A_n denote the event that $A_{n,m}$ holds for all $m \in \{1, 2, \ldots, M\}$, i.e., $A_n = \bigcap_{m=1}^{M} A_{n,m}$.

Lemma 71 *If* $\gamma_{\mathrm{RRT}^*} > 4 \left(\frac{\mu(\mathcal{X}_{\mathrm{free}})}{\zeta_d}\right)^{1/d}$*, then* A_n *occurs for all large* n*, with probability one, i.e.,*

$$\mathbb{P}\left(\liminf_{n\to\infty} A_n\right) = 1.$$

Proof The proof of this result is based on a Poissonization argument. Let Poisson(λ) be a Poisson random variable with parameter $\lambda = \theta\, n$, where $\theta \in (0, 1)$ is a constant independent of n. Consider the point process that consists of exactly Poisson$(\theta\, n)$ points, i.e., $\{X_1, X_2, \ldots, X_{\mathrm{Poisson}(\theta\, n)}\}$. This point process is a Poisson point process with intensity $\theta\, n\, /\mu(X_{\mathrm{free}})$ by Lemma 11.

Let $\tilde{A}_{n,m}$ denote the event that there exists two vertices X_i and $X_{i'}$ in the vertex set of the RRT* algorithm such that X_i and $X_{i'}$ are connected with an edge in $\tilde{G}_n$, where $\tilde{G}_n$ is the graph returned by the RRT* when the algorithm is run for Poisson$(\theta\, n)$ many iterations, i.e., Poisson$(\theta\, n)$ samples are drawn from $\mathcal{X}_{\mathrm{free}}$.

Clearly, $\mathbb{P}(A^c_{n,m}) = \mathbb{P}(\tilde{A}^c_{n,m} \,|\, \{\mathrm{Poisson}(\theta\, n) = n\})$. Moreover,

$$\mathbb{P}(A^c_{n,m}) \;\le\; \mathbb{P}(\tilde{A}^c_{n,m}) + \mathbb{P}(\{\mathrm{Poisson}(\theta\, n) > n\}).$$

since $\mathbb{P}(A^c_{n,m})$ is non-increasing with n (see, e.g., Penrose, 2003). Since $\theta < 1$, $\mathbb{P}(\{\mathrm{Poisson}(\theta\, n) > n\}) \;\le\; e^{-a\, n}$, where $a > 0$ is a constant independent of n.

To compute $\mathbb{P}(\widetilde{A}^c_{n,m})$, a number of definitions are provided. Let $N_{n,m}$ denote the number of vertices that lie in the interior of $B_{n,m}$. Clearly, $\mathbb{E}[N_{n,m}] = \frac{\zeta_d\, \gamma^d_{\mathrm{RRT}^*}}{\mu(X_{\mathrm{free}})} \log n$, for all $m \in \{1, 2, \ldots, M_n\}$. For notational simplicity, define $\alpha := \frac{\zeta_d\, \gamma^d_{\mathrm{RRT}^*}}{\mu(X_{\mathrm{free}})}$. Let $\epsilon \in (0,1)$ be a constant independent of n. Define the event

$$C_{n,m,\epsilon} \quad := \quad \{N_{n,m} \geq (1-\epsilon)\, \mathbb{E}[N_{n,m}]\} = \{N_{n,m} \geq (1-\epsilon)\, \alpha\, \log n\}$$

Since $N_{n,m,\epsilon}$ is binomially distributed, its large deviations from its mean can be bounded as follows (Penrose, 2003),

$$\mathbb{P}\left(C^c_{n,m,\epsilon}\right) = \mathbb{P}(\{N_{n,m,\epsilon} \leq (1-\epsilon)\, \mathbb{E}[N_{n,m}]\}) \leq e^{-\alpha\, H(\epsilon)\, \log n} = n^{-\alpha H(\epsilon)},$$

where $H(\epsilon) = \epsilon + (1-\epsilon)\log(1-\epsilon)$. Notice that $H(\epsilon)$ is a continuous function of ϵ with $H(0) = 0$ and $H(1) = 1$. Hence, $H(\epsilon)$ can be made arbitrary close to one by taking ϵ close to one.

Then,

$$\begin{aligned}\mathbb{P}(\tilde{A}^c_{n,m}) \quad &= \quad \mathbb{P}(\tilde{A}^c_{n,m} \,|\, C_{n,m,\epsilon} \cap C_{n,m+1,\epsilon})\, \mathbb{P}(C_{n,m,\epsilon} \cap C_{n,m+1,\epsilon}) \\ &\quad\quad +\mathbb{P}(\tilde{A}^c_{n,m} \,|\, (C_{n,m,\epsilon} \cap C_{n,m+1,\epsilon})^c)\, \mathbb{P}((C_{n,m,\epsilon} \cap C_{n,m+1,\epsilon})^c) \\ &\leq \quad \mathbb{P}(\tilde{A}^c_{n,m} \,|\, C_{n,m,\epsilon} \cap C_{n,m+1,\epsilon})\, \mathbb{P}(C_{n,m,\epsilon} \cap C_{n,m+1,\epsilon}) + \mathbb{P}(C^c_{n,m,\epsilon}) + \mathbb{P}(C^c_{n,m+1,\epsilon}),\end{aligned}$$

where the last inequality follows from the union bound.

First, using the spatial independence of the underlying point process,

$$\mathbb{P}\left(C_{n,m,\epsilon} \cap C_{n,m+1,\epsilon}\right) = \mathbb{P}\left(C_{n,m,\epsilon}\right)\, \mathbb{P}\left(C_{n,m+1,\epsilon}\right) \leq n^{-2\,\alpha\, H(\epsilon)}.$$

Second, observe that $\mathbb{P}(A^c_{n,m} \,|\, N_{n,m} = k, N_{n,m+1} = k')$ is a non-increasing function of both k and k', since the probability of the event $\tilde{A}_{n,m}$ can not increase with the increasing number of points in both balls, $B_{n,m}$ and $B_{n,m+1}$. Then,

$$\begin{aligned}\mathbb{P}(\tilde{A}^c_{n,m} \,|\, C_{n,m,\epsilon} \cap C_{n,m+1,\epsilon}) \quad &= \quad \mathbb{P}(\tilde{A}^c_{n,m} \,|\, \{N_{n,m} \geq (1-\epsilon)\, \alpha\, \log N_{n,m}, N_{n,m+1} \geq (1-\epsilon)\, \alpha\, \log N_{n,m+1}\}) \\ &\leq \quad \mathbb{P}(\tilde{A}^c_{n,m} \,|\, \{N_{n,m} = (1-\epsilon)\, \alpha\, \log N_{n,m}, N_{n,m+1} = (1-\epsilon)\, \alpha\, \log N_{n,m+1}\})\end{aligned}$$

The term on the right hand side is one minus the probability that the maximum of $\alpha\, \log n$ number of uniform samples drawn from $[0,1]$ is smaller than the minimum of $\alpha\, \log n$ number of samples again drawn from $[0,1]$, where all the samples are drawn independently. This probability can be calculated as follows. From the order statistics of uniform distribution, the minimum of $\alpha\, \log n$ points sampled independently and uniformly from $[0,1]$ has the following probability distribution function:

$$f_{\min}(x) = \frac{(1-x)^{\alpha\, \log n - 1}}{\mathtt{Beta}(1, \alpha\, \log(n))},$$

where $\mathtt{Beta}(\cdot,\cdot)$ is the Beta function (also called the Euler integral) (Abramowitz and Stegun, 1964). The maximum of the same number of independent uniformly distributed random variables with support $[0,1]$ has the following cumulative distribution function:

$$F_{\max}(x) = x^{\alpha \log n}$$

Then,

$$\begin{aligned}
\mathbb{P}(\tilde{A}^c_{n,m}\,|\,C_{n,m,\epsilon}\cap C_{n,m+1,\epsilon}) &\le \int_0^1 F_{\max}(x)\,f_{\min}(x)\,dx \\
&= \frac{\texttt{Gamma}((1-\epsilon)\,\alpha\,\log n)\,\texttt{Gamma}((1-\epsilon)\,\epsilon\,\log n)}{2\,\texttt{Gamma}(2(1-\epsilon)\,\alpha\,\log(n))} \\
&\le \frac{((1-\epsilon)\,\alpha\,\log n)!\,((1-\epsilon)\,\alpha\,\log n)!}{2\,(2\,(1-\epsilon)\,\alpha\,\log n)!} \\
&= \frac{((1-\epsilon)\,\alpha\,\log n)!}{2(2(1-\epsilon)\,\alpha\,\log n)(2(1-\epsilon)\,\alpha\,\log n-1)\cdots 1} \\
&\le \frac{1}{2^{(1-\epsilon)\,\alpha\,\log n}} = n^{-\,\log(2)\,(1-\epsilon)\,\alpha},
\end{aligned}$$

where $\texttt{Gamma}(\cdot)$ is the gamma function (Abramowitz and Stegun, 1964).

Then,

$$\mathbb{P}(\tilde{A}^c_{n,m}) \le n^{-\alpha\left(2\,H(\epsilon)+\log(2)\,(1-\epsilon)\right)} + 2\,n^{-\alpha\,H(\epsilon)}.$$

Since $2\,H(\epsilon)+\log(2)\,(1-\epsilon)$ and $H(\epsilon)$ are both continuous and increasing in the interval $(0.5,1)$, the former is equal to $2-\log(4) > 0.5$ and the latter is equal to 1 as ϵ approaches one from below, there exists some $\bar{\epsilon}\in(0.5,1)$ such that both $2\,H(\bar{\epsilon})+\log(2)\,(1-\bar{\epsilon}) > 0.5$ and $H(\bar{\epsilon}) > 0.5$. Thus,

$$\mathbb{P}(\tilde{A}^c_{n,m}) \le n^{-\alpha/2} + 2\,n^{-\alpha/2} = 3\,n^{-\alpha/2}.$$

Hence,

$$\begin{aligned}
\mathbb{P}(A^c_{n,m}) &\le \mathbb{P}(\tilde{A}^c_{n,m}) + \mathbb{P}(\text{Poisson}(\theta\,n) > n) \\
&\le 3\,n^{-\alpha/2} + e^{-a\,n}
\end{aligned}$$

Recall that A_n denotes the event that $A_{n,m}$ holds for all $m\in\{1,2,\ldots,M_n\}$. Then,

$$\mathbb{P}(A^c_n) \;=\; \mathbb{P}\left(\left(\bigcap\nolimits_{m=1}^{M_n} A_{n,m}\right)^c\right) \;=\; \mathbb{P}\left(\bigcup\nolimits_{m=1}^{M_n} A^c_{n,m}\right) \;\le\; \sum_{m=1}^{M_n}\mathbb{P}\left(A^c_{n,m}\right) \;=\; M_n\,\mathbb{P}(A^c_{n,1}),$$

where the last inequality follows from the union bound. The number of balls in B_n can be bounded as

$$|B_n| \;=\; M_n \;\le\; \beta\left(\frac{n}{\log n}\right)^{1/d},$$

where β is a constant. Combining this with the inequality above,

$$\mathbb{P}(A^c_n) \;\le\; \beta\left(\frac{n}{\log n}\right)^{1/d}\left(3\,n^{-\alpha/2}+e^{-a\,n}\right),$$

which is summable for $\alpha > 2\,(1+1/d)$. Thus, by the Borel-Cantelli lemma, the probability that A^c_n occurs infinitely often is zero, i.e., $\mathbb{P}(\limsup_{n\to\infty} A^c_n) = 0$, which implies that A_n occurs for all large n with probability one, i.e., $\mathbb{P}(\liminf_{n\to\infty} A_n) = 1$. ☐

G.4 Convergence to the optimal path

The proof of the following lemma is similar to that of Lemma 55, and is omitted here.

Let P_n denote the set of all paths in the graph returned by RRT* algorithm at the end of n iterations. Let σ'_n be the path that is closest to σ_n in terms of the bounded variation norm among all those paths in P_n, i.e., $\sigma'_n := \min_{\sigma' \in P_n} \|\sigma' - \sigma_n\|$.

Lemma 72 *The random variable* $\|\sigma'_n - \sigma_n\|_{\mathrm{BV}}$ *converges to zero almost surely, i.e.,*

$$\mathbb{P}\left(\left\{\lim_{n\to\infty} \|\sigma'_n - \sigma_n\|_{\mathrm{BV}} = 0\right\}\right) = 1.$$

A corollary of the lemma above is that $\lim_{n\to\infty} \sigma'_n = \sigma^*$ with probability one. Then, the result follows by the robustness of the optimal solution (see the proof of Lemma 56 for details).